

# On Demand Operating Environment: Managing the Infrastructure (Virtualization Engine Update)

Introduction to the On Demand  
Operating Environment

Automate and virtualize the IT  
environment

Start deploying today



Bart Jacob  
Surrey Mui  
Jatinder Pannu  
Sungsim Park  
Hugues Raguet  
Jack Schneider  
Laurent Vanel

Christian Matthys  
Paola Bari  
Emmanuel Lieurain  
Dominique Salomon  
Lynn Winkelbauer





International Technical Support Organization

**On Demand Operating Environment:  
Managing the Infrastructure  
(Virtualization Engine Update)**

June 2005

**Note:** Before using this information and the product it supports, read the information in “Notices” on page xiii.

## **Second Edition (June 2005)**

This edition adds information about the IBM Virtualization Engine; this information applies to version 1, release 1 of the IBM Virtualization Engine

**© Copyright International Business Machines Corporation 2004, 2005. All rights reserved.**

Note to U.S. Government Users Restricted Rights -- Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

# Contents

<b>Figures</b> .....	xi
<b>Notices</b> .....	xiii
Trademarks .....	xiv
<b>Preface</b> .....	xv
The team that wrote this redbook .....	xvi
Become a published author .....	xx
Comments welcome .....	xx
<b>Summary of changes</b> .....	xxiii
June 2005, Second Edition .....	xxiii
<b>Part 1. Overview</b> .....	1
<b>Chapter 1. Introduction</b> .....	3
1.1 Getting to On Demand Business .....	4
1.2 Infrastructure to support an On Demand Business .....	6
1.3 Capabilities .....	7
1.3.1 Integration capabilities .....	8
1.3.2 Infrastructure management capabilities .....	9
1.4 On Demand Operating Environment architecture .....	10
1.5 Summary .....	12
<b>Chapter 2. Infrastructure management overview</b> .....	13
2.1 Business drivers .....	14
2.2 Framework for infrastructure management .....	14
2.3 Automation .....	17
2.3.1 Business Service Management .....	18
2.3.2 Policy-based Orchestration .....	19
2.3.3 Availability .....	20
2.3.4 Security .....	21
2.3.5 Optimization .....	22
2.3.6 Provisioning .....	24
2.4 Virtualization .....	25
2.4.1 The value of virtualization .....	26
2.4.2 Server virtualization .....	27
2.4.3 Storage virtualization .....	33
2.4.4 Network virtualization .....	37

2.4.5 Distributed systems . . . . .	39
2.4.6 The IBM Virtualization Engine . . . . .	43
<b>Chapter 3. The IBM Virtualization Engine . . . . .</b>	<b>45</b>
3.1 Overview of the IBM Virtualization Engine . . . . .	46
3.2 Systems technologies . . . . .	47
3.2.1 Systems technologies for the zSeries family . . . . .	47
3.2.2 Systems technologies integrated into the pSeries . . . . .	49
3.2.3 Systems technologies integrated into the iSeries . . . . .	50
3.2.4 Systems technologies integrated into xSeries and BladeCenter . . . . .	51
3.3 Systems services . . . . .	52
3.3.1 Suite for Servers . . . . .	52
3.3.2 Suite for Storage . . . . .	61
3.3.3 System services summary . . . . .	63
3.4 Summary . . . . .	65
<b>Part 2. How to's for managing the Infrastructure . . . . .</b>	<b>67</b>
<b>Chapter 4. How to secure access and control of information, resources, and applications . . . . .</b>	<b>69</b>
4.1 Vision . . . . .	70
4.1.1 Web Services security components . . . . .	70
4.1.2 Hardware and software security mechanisms . . . . .	70
4.1.3 Glimpse of the future . . . . .	72
4.2 How to get started today . . . . .	72
<b>Chapter 5. How to provide scalable and consistent management and control of operations for end-to-end business systems . . . . .</b>	<b>73</b>
5.1 Vision . . . . .	74
5.1.1 The role of standards . . . . .	74
5.1.2 Orchestration and provisioning . . . . .	76
5.1.3 The Service Oriented Architecture model . . . . .	76
5.1.4 Automation . . . . .	77
5.2 How to get started today . . . . .	78
5.2.1 Enterprise Workload Management component . . . . .	78
5.2.2 IBM Tivoli Business Systems Manager component . . . . .	79
5.2.3 Virtualization Engine Console component . . . . .	79
5.2.4 IBM Director Multiplatform component . . . . .	80
5.2.5 IBM Tivoli Enterprise Console component . . . . .	80
5.2.6 IBM Tivoli Monitoring components . . . . .	81
5.2.7 IBM Tivoli Intelligent ThinkDynamic Orchestrator component . . . . .	81
5.2.8 The storage components . . . . .	81
5.2.9 The Electronic Service Agent component . . . . .	82
5.2.10 The automation environment . . . . .	83

5.3 Key products to start with .....	85
<b>Chapter 6. How to avoid system failures and take automated action to resolve problems .....</b>	<b>87</b>
6.1 Vision.....	88
6.2 How to get started today .....	88
6.2.1 Hardware solutions .....	88
6.2.2 Software and services solutions .....	90
6.3 Key products to start with .....	96
<b>Chapter 7. How to protect systems from intrusions and threats using monitor and alert systems .....</b>	<b>97</b>
7.1 Vision.....	98
7.2 How to get started today .....	98
7.2.1 The IBM Integrated Security Solution for Cisco Networks .....	98
7.2.2 The IBM Tivoli Security Management environment .....	101
7.2.3 z/OS environment solutions .....	104
7.3 Key products to start with .....	105
<b>Chapter 8. How to monitor systems to allow establishment of business SLAs and automate detection and remediation of violations</b>	<b>107</b>
8.1 Vision.....	108
8.2 How to get started today .....	109
8.2.1 The IBM Tivoli Business Service Management solution .....	109
8.3 Key products to start with .....	112
<b>Chapter 9. How to reduce the time and cost to re-purpose IT resources to meet business requirements .....</b>	<b>113</b>
9.1 Vision.....	114
9.1.1 Architecture approach: the MAPE Loop .....	114
9.1.2 Optimizing existing resources .....	115
9.1.3 Storage approach .....	116
9.2 How to get started today .....	116
9.2.1 The MAPE Loop implementation .....	116
9.2.2 The role of the grid .....	120
9.2.3 Storage implementations .....	120
9.3 Key products to start with .....	123
<b>Chapter 10. How to map IT resources used by various business processes of an end-to-end solution .....</b>	<b>125</b>
10.1 Vision.....	126
10.2 How to get started today .....	126
10.2.1 Understand the usage of data.....	126
10.2.2 Understand the usage of specific components .....	127

10.2.3 Sample scenario . . . . .	128
10.3 Key products to start with . . . . .	130
<b>Chapter 11. How to consolidate and simplify the IT infrastructure . . . .</b>	<b>131</b>
11.1 Vision. . . . .	135
11.2 How to get started today . . . . .	136
11.2.1 zSeries implementation example . . . . .	137
11.2.2 Storage implementation example . . . . .	138
11.2.3 BladeCenter implementation example . . . . .	140
11.3 Key products to start with . . . . .	140
<b>Chapter 12. How to optimize utilization and pool resources across a heterogeneous environment . . . . .</b>	<b>143</b>
12.1 Vision. . . . .	144
12.1.1 Open Grid Services Architecture. . . . .	144
12.1.2 Workload management . . . . .	145
12.2 How to get started today . . . . .	145
12.2.1 The zSeries example . . . . .	145
12.2.2 Examples of LPAR . . . . .	146
12.2.3 The grid benefit . . . . .	146
12.2.4 Mixing partitioning and grid capabilities . . . . .	149
12.2.5 Infrastructure management tools . . . . .	149
12.2.6 Storage focus . . . . .	152
12.3 Key products to start with . . . . .	155
<b>Chapter 13. How to provision system resources in order to meet business demands . . . . .</b>	<b>157</b>
13.1 Vision. . . . .	158
13.2 How to get started today . . . . .	158
<b>Chapter 14. How to monitor end-to-end applications, their topology, and their resources . . . . .</b>	<b>159</b>
14.1 Vision. . . . .	160
14.2 How to get started today . . . . .	161
14.2.1 IBM Tivoli Monitoring for Transaction Performance . . . . .	161
14.2.2 Enterprise Workload Manager . . . . .	162
14.3 Key products to start with . . . . .	163
<b>Part 3. Infrastructure management: Detailed scenarios . . . . .</b>	<b>165</b>
<b>Chapter 15. How to secure access and control of information, resources, and applications . . . . .</b>	<b>167</b>
15.1 Introduction . . . . .	168
15.2 General strategy . . . . .	169



15.3	Solution components . . . . .	170
15.3.1	Identity management . . . . .	170
15.3.2	Privacy management . . . . .	171
15.3.3	Security management console . . . . .	171
15.3.4	Data protection . . . . .	172
15.4	Scenario . . . . .	172
15.4.1	Business context . . . . .	172
15.4.2	Current environment . . . . .	173
15.4.3	Business objectives . . . . .	175
15.4.4	Technical objectives . . . . .	175
15.4.5	Solution approach . . . . .	181
15.5	Product positioning . . . . .	198
15.5.1	Identity management . . . . .	199
15.5.2	Privacy Control Management . . . . .	203
15.6	Linkages . . . . .	204
15.7	Glimpse of the future . . . . .	212
15.8	Summary . . . . .	213
 <b>Chapter 16. How to provision system resources according to business demands . . . . .</b>		<b>215</b>
16.1	Introduction . . . . .	216
16.2	General strategy . . . . .	216
16.3	Solution components . . . . .	217
16.4	Scenario . . . . .	218
16.4.1	Current environment . . . . .	219
16.4.2	Business objectives . . . . .	221
16.4.3	Technical objectives . . . . .	222
16.4.4	Solution approach . . . . .	222
16.4.5	Benefits and summary . . . . .	240
16.5	Product positioning . . . . .	241
16.5.1	Self-configuring IBM @server . . . . .	241
16.5.2	Self-configuring IBM TotalStorage . . . . .	246
16.5.3	Hypervisors and workload managers . . . . .	246
16.5.4	IBM Tivoli Provisioning Manager . . . . .	250
16.5.5	IBM Tivoli Intelligent ThinkDynamic Orchestrator . . . . .	250
16.6	Linkages . . . . .	251
16.7	Glimpse of the future . . . . .	253
16.8	Summary . . . . .	254
 <b>Chapter 17. How to balance workloads in the network . . . . .</b>		<b>255</b>
17.1	Introduction . . . . .	256
17.1.1	Overview of the CISCO CSM . . . . .	257
17.1.2	Overview of EWLM's load balancing recommendations . . . . .	258

17.1.3 The Server/Application State Protocol (SASP) . . . . .	258
17.2 Configuring the components . . . . .	259
17.2.1 Configuring the Catalyst 6509 and the CSM for EWLM . . . . .	260
17.2.2 Configuring EWLM for load balancing. . . . .	261
17.2.3 Failover considerations . . . . .	262
17.3 Network balancing example . . . . .	262
17.3.1 Network and application topology . . . . .	263
17.3.2 Load balancing configurations . . . . .	263
17.4 Lessons learned . . . . .	267
17.4.1 Best practices . . . . .	267
17.4.2 Special EWLM benefits to load balancing. . . . .	268
17.4.3 Troubleshooting . . . . .	268
17.5 Conclusion. . . . .	269
<b>Chapter 18. How to consolidate, simplify, and optimize the storage IT infrastructure. . . . .</b>	<b>271</b>
18.1 Introduction . . . . .	272
18.2 General strategy . . . . .	273
18.3 Solution components. . . . .	274
18.4 Scenario . . . . .	274
18.4.1 Current environment . . . . .	274
18.4.2 Business objectives. . . . .	276
18.4.3 Technical objectives . . . . .	277
18.4.4 Solution approach . . . . .	277
18.4.5 Benefits and summary . . . . .	291
18.5 Product positioning: Conclusion . . . . .	292
18.5.1 IBM TotalStorage SAN Volume Controller . . . . .	292
18.5.2 IBM TotalStorage SAN File System . . . . .	293
18.5.3 IBM TotalStorage Productivity Center. . . . .	293
18.6 Summary . . . . .	293
<b>Chapter 19. How to monitor using EWLM. . . . .</b>	<b>295</b>
19.1 Introduction . . . . .	296
19.2 General strategy . . . . .	296
19.3 Solution components. . . . .	297
19.4 Scenario . . . . .	298
19.4.1 Current environment . . . . .	298
19.4.2 Business objective . . . . .	299
19.4.3 Technical objectives . . . . .	299
19.5 Scenario implementation. . . . .	299
19.5.1 Architecture components . . . . .	300
19.5.2 Implementation of IBM EWLM . . . . .	302
19.5.3 Monitoring scenario. . . . .	315

19.5.4	How to use EWLM to debug/diagnose a performance problem . .	320
19.5.5	Benefits and summary . . . . .	321
19.6	Glimpse of the future . . . . .	322
19.7	Summary . . . . .	322
<b>Part 4.</b>	<b>Appendixes . . . . .</b>	<b>323</b>
<b>Appendix A.</b>	<b>Getting Started with the Virtualization Engine . . . . .</b>	<b>325</b>
A.1	“Start Your Engines” Workshops . . . . .	326
A.2	VE Consulting/Assessment study engagement overview . . . . .	328
A.2.1	Objectives . . . . .	328
A.2.2	Offerings . . . . .	329
A.2.3	Methodology . . . . .	330
A.3	The Planning Advisor . . . . .	331
<b>Appendix B.</b>	<b>Standards overview . . . . .</b>	<b>333</b>
B.1	Open source . . . . .	334
B.2	Standards organizations . . . . .	334
B.2.1	IETF - Internet Engineering Task Force . . . . .	334
B.2.2	W3C - World Wide Web Consortium . . . . .	335
B.2.3	JCP - Java Community Process . . . . .	335
B.2.4	OASIS - Organization for the Advancement of Structured Information Standards . . . . .	336
B.2.5	WS-I - Web Services Interoperability Organization . . . . .	336
B.2.6	DMTF - Distributed Management Task Force . . . . .	337
B.2.7	GGF - Global Grid Forum . . . . .	337
B.2.8	OMG - Object Management Group . . . . .	338
B.3	Key standards . . . . .	338
B.3.1	XML standards . . . . .	339
B.3.2	SOAP - Simple Object Access Protocol . . . . .	341
B.3.3	WSDL - Web Services Description Language . . . . .	342
B.3.4	UDDI - Universal Description, Discovery, and Integration . . . . .	342
B.3.5	WS-I Basic Profile 1.0a . . . . .	342
B.3.6	WS-Security - Web Services Security . . . . .	343
B.3.7	OGSA - Open Grid Services Architecture . . . . .	343
B.3.8	OGSI - Open Grid Services Infrastructure . . . . .	344
B.3.9	UML - Unified Modeling Language . . . . .	345
B.3.10	MDA - Model Driven Architecture . . . . .	346
B.3.11	CIM - Common Information Model . . . . .	346
B.3.12	Open Group . . . . .	347
<b>Related publications</b>	<b>. . . . .</b>	<b>349</b>
IBM Redbooks	. . . . .	349
Other publications	. . . . .	350

Online resources ..... 351

How to get IBM Redbooks ..... 354

Help from IBM ..... 354

**Index** ..... 355

# Figures

1-1	Positioning the On Demand Operating Environment . . . . .	4
1-2	The On Demand Operating Environment . . . . .	8
1-3	On Demand Operating Environment architecture . . . . .	11
2-1	Customer Infrastructure progression . . . . .	15
2-2	Capabilities for IT simplification . . . . .	16
2-3	Types of server partitioning . . . . .	28
2-4	Storage Virtualization components . . . . .	34
2-5	Web Services Resource Framework (WS-RF) . . . . .	42
3-1	Enterprise workload monitoring . . . . .	53
3-2	The Virtualization Engine console . . . . .	58
3-3	Virtualization Engine integration . . . . .	64
3-4	Stages of virtualization solutions . . . . .	66
6-1	Electronic service agent architecture . . . . .	94
6-2	IBM Tivoli Monitoring high-level architecture . . . . .	95
7-1	IBM Tivoli Security Management blueprint . . . . .	102
8-1	Business Service Management products integration . . . . .	110
9-1	The Autonomic Computing MAPE Loop . . . . .	115
11-1	Types of consolidation . . . . .	133
14-1	Automation Blueprint . . . . .	160
15-1	ITSO-Electronics.com data network . . . . .	174
15-2	Problem with provisioning new user . . . . .	177
15-3	Problem when managing user accounts . . . . .	177
15-4	Problem of de-provisioning users . . . . .	178
15-5	Problem with new initiatives . . . . .	179
15-6	Current ITSO-Electronics.com architecture . . . . .	179
15-7	Capabilities of an identity management solution . . . . .	181
15-8	IBM's Integrated Identity Management solution . . . . .	182
15-9	Initial IBM Identity Management architecture . . . . .	185
15-10	Security architecture with internal WebSEAL . . . . .	187
15-11	Security architecture with high availability, scalability, and resilience .	190
15-12	New internal and highly sensitive Web application . . . . .	193
15-13	Shifting value of identity management . . . . .	197
15-14	Identity management blueprint . . . . .	199
15-15	Product linkage overview . . . . .	206
15-16	User provisioning workflow . . . . .	207
15-17	User provisioning . . . . .	208
15-18	Access control and single sign-on to various back end applications .	209
16-1	Current Stocks Online and Employee Online applications . . . . .	219
16-2	Orchestrated provisioning in San Diego data center . . . . .	225

16-3	IBM Tivoli Intelligent ThinkDynamic Orchestrator architecture . . . . .	227
16-4	IBM Web Infrastructure Orchestration for ITSO-Electronics.com . . . .	229
16-5	Logical network data of ITSO-Electronics.com San Diego data center	231
16-6	High-level asset relationship in ITSO-Electronics.com . . . . .	233
16-7	Additional resource is required for Stocks Online application . . . . .	236
16-8	Removal of resource for Stocks Online application . . . . .	237
16-9	Provisioning of server to the Employee Online application . . . . .	239
16-10	VMware ESX architecture . . . . .	247
16-11	Use of resources on a zSeries. . . . .	248
16-12	iSeries IT optimization . . . . .	249
16-13	Linkage of provisioning products . . . . .	253
17-1	CSM and EWLM interaction. . . . .	257
17-2	Typical network and application topology . . . . .	263
17-3	EWLM Control Center . . . . .	264
17-4	System topology view . . . . .	264
18-1	Current Customer Storage configuration . . . . .	275
18-2	New IT Storage Architecture . . . . .	279
18-3	SVC and SFS detailed configuration . . . . .	284
18-4	Visualization of the storage topology . . . . .	285
18-5	Storage identification . . . . .	286
19-1	Web applications layout . . . . .	298
19-2	EWLM architecture . . . . .	300
19-3	Exception monitor . . . . .	315
19-4	Service class details . . . . .	316
19-5	Server topology for the Response time service class . . . . .	317
19-6	Server topology data . . . . .	318
19-7	Managed server details report . . . . .	320
A-1	Possible tasks before implementation . . . . .	326
A-2	Start Your Engines lab environment . . . . .	327
A-3	Virtualization Engine Assessment offerings. . . . .	329
B-1	Key standards . . . . .	339
B-2	Grid specifications and related standards . . . . .	345

# Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

*IBM Director of Licensing, IBM Corporation, North Castle Drive Armonk, NY 10504-1785 U.S.A.*

*The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:* INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.



This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

## COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrates programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. You may copy, modify, and distribute these sample programs in any form without payment to IBM for the purposes of developing, using, marketing, or distributing application programs conforming to IBM's application programming interfaces.

# Trademarks

The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

AIX®	HiperSockets™	pSeries®
AIX 5L™	HyperSwap™	RACF®
BladeCenter™	Hypervisor™	Redbooks™
CICS®	IBM®	Redbooks (logo)  ™
Cloudscape™	ibm.com®	RMF™
Component Business Model™	IMS™	RS/6000®
DB2 Universal Database™	Informix®	S/390®
DB2®	iSeries™	SecureWay®
DFSMSHsm™	i5/OS™	Storage Tank™
Domino®	Lotus®	Tivoli®
e-business on demand™	Micro-Partitioning™	Tivoli Enterprise™
ECKD™	MVS™	Tivoli Enterprise Console®
Electronic Service Agent™	NetView®	TotalStorage®
Enterprise Storage Server®	OS/390®	Virtualization Engine™
@server®	OS/400®	VSE/ESA™
 server®	Parallel Sysplex®	WebSphere®
FlashCopy®	POWER™	xSeries®
GDPS®	POWER4™	z/OS®
Geographically Dispersed	POWER5™	z/VM®
Parallel Sysplex™	PowerPC®	zSeries®
HACMP™	PR/SM™	

The following terms are trademarks of other companies:

Intel, Intel Inside (logos), MMX, Centrino and Pentium are trademarks of Intel Corporation in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Other company, product, and service names may be trademarks or service marks of others.



# Preface

This redbook (along with its companion volume, *On Demand Operating Environment: Creating Business Flexibility*, SG24-6633), provides an insight into the kind of operating environment required to support an On Demand Business.

It provides an overview of the architecture of an On Demand Operating Environment and describes in detail the components that are required to manage the infrastructure. To meet the business needs of being responsive, variable, focused, and resilient, an On Demand Operating Environment must be integrated, autonomic, virtualized, and open. Though these attributes are all interrelated, this redbook focuses on the automation and virtualization components as they enable efficient infrastructure management.

This second edition adds more detailed information about the IBM® Virtualization Engine™, which is a key feature related to the virtualization aspects for the infrastructure, and which was not available when the first edition was written.

A complete On Demand Operating Environment is a vision and a goal that many enterprises aspire to reach. However, it is not something that will be attained overnight or by installing a specific set of products. It is something that will be reached through a step-wise progression.

This redbook provides descriptions of several approaches that one can choose to start implementing pieces of an On Demand Operating Environment today. Which approach is right for the reader will depend on their specific business environment and their immediate needs.

The book is organized into three main parts:

1. Part 1, "Overview" on page 1, introduces the On Demand Operating Environment and describes how its two main principles, the automation and the virtualization, help to simplify an IT infrastructure management.
2. Part 2, "How to's for managing the Infrastructure" on page 67, introduces some specific aspects of the management of an IT infrastructure, describing for each of them the IBM vision and the role of specific components in the simplification process.
3. Part 3, "Infrastructure management: Detailed scenarios" on page 165, details very specific scenarios, explaining how to implement the components of a solution and demonstrating the value of the solution.

Our objective is to help the reader better understand what an On Demand Operating Environment is and how they can take steps today to start putting one in place. This redbook does not go into detailed implementation plans for each technology or product it references, but rather provides a level of information sufficient for the reader to start building a strategy and architecture best suited for their needs. Product-specific details can be obtained from product documentation and product-related redbooks.

## The team that wrote this redbook

The second edition of this redbook was produced by a team of specialists working partly at the International Technical Support Organization (ITSO), Poughkeepsie Center, New York, and partly at the Products and Solutions Support Center (PSSC), which is part of the EMEA Advanced Technical Support (ATS), located in Montpellier, France:

**Christian Matthys** spent more than 20 years at IBM as a System Engineer, working with large mainframe-oriented customers. He spent three years as an ITSO project leader on assignment in Poughkeepsie, NY. In 2000 he joined the EMEA Design Center for On Demand Business, working with customers' projects to make use of the leading edge technologies, in particular the Autonomic Computing technologies. He was part of the launching team of the Start Your Engines workshops to promote the IBM Virtualization Engine. He works as an ITSO project leader based in the PSSC. He is a certified Consulting IBM IT Specialist.

**Paola Bari** is an Advisory Programmer at the International Technical Support Organization, Poughkeepsie Center. She has 22 years of experience as a systems programmer in OS/390®, z/OS®, and Parallel Sysplex®, including several years of experience in WebSphere® MQ and WebSphere Application Server.

**Emmanuel Lieurain** engaged in studies of Artificial Intelligence and Multi-Agents Systems in 1998. He holds a Master's degree gained by working on distributed simulation written in Java™. Over 2 years, he has taught many courses for IBM Learning Services in a Business Partner company. He started working for IBM in 2001, as a Java developer and a WebSphere Specialist. As a member of the Design Center for On Demand Business, he participates in architecture design sessions, proof of concepts, and porting. He is also involved in pSeries® and multiplatform benchmarks for tuning and optimizing WebSphere Environments. In 2004, Emmanuel was involved in the Autonomic Computing initiative. He is now promoting the IBM Virtualization Engine through technical workshops in a program named Start Your Engines, of which he is the EMEA technical leader.

**Dominique Salomon** has been with IBM for 6 years and has 15 years of experience on UNIX, High Availability, and Storage domains. He is currently part of the PSSC TotalStorage® team as an education leader. He has specialized in AIX® and storage virtualization products such as SAN Volume Controller (SVC) and SAN File System (SFS). Dominique has also a strong background in complex High Availability and Storage solutions industrialization processes.


**Lynn Winkelbauer** is a System Engineer for IBM, Poughkeepsie. She has 21 years of experience in the zSeries® area, working in performance, Parallel Sysplex, and e-business on demand™. She holds a Bachelor of Science degree in Computer Science and Business Administration. Her areas of expertise include Linux for zSeries, CICS® Transaction Server, WebSphere Portal Server, IBM Virtualization Engine and performance analysis.

Thanks to the following individuals from across IBM who assisted with the development of the second edition of this redbook:

Don Bagwell	Patrick Kappeler
Katalin Bartfai-Walcott	Yvonne Lyon
Boas Betzler	Tim McCrimmon
Alan Bivens	Annette Miller
Mark Cathcart	Tom Monza
Werner Ederer	Maddie Nick
David Goodman	Kristi Schultz
Utte Gindl	Julie Shore
Heather Hinton	Joe Winkelbauer
Bart Jacob	Jim Xiong

Our apologies to any others that may have been inadvertently left off this list. Your help was appreciated.

The first version of this redbook was produced by a team of specialists from around the world working at the International Technical Support Organization, Austin Center.



*From left to right: Bart Jacob, Laurent Vanel, Hugues Raguet, Sungsim Park, Surrey Mui, Jatinder Pannu, Jack Schneider*

**Bart Jacob** is a Senior Consulting IT Specialist in the IBM International Technical Support Organization, Austin Center. He has 23 years of experience providing technical support across a variety of IBM products and technologies, including communications, object-oriented software development, and systems management. He has over 10 years of experience at the ITSO, where he has been writing IBM Redbooks™ and creating and teaching workshops around the world on a variety of topics. He holds a Master's degree in Numerical Analysis from Syracuse University.

**Surrey Mui** is an Advisory IT Specialist at IBM Global Services, Hong Kong. She has over 10 years of experience in the Information Technology industry. Before joining IGS-HK in August 2003, she was a member of the Advanced Customer Engineering team working in the Tivoli® Security Business Unit in Australia (part of the IBM Software Group). Her area of expertise is in providing architecture and integration solutions for customers using the Tivoli Access Manager product suite. This includes building specialized development modules for customers based on the Access Manager product suite. Surrey was a member of the DASCOT team when they were acquired by IBM. Surrey holds a degree in Business Computing from Griffith University, Australia.

**Jatinder Pannu** is a Senior Consulting Client IT Architect in IBM US. He has over twenty years of experience in the IT industry, and has worked for IBM for nineteen years planning requirements and designing systems with customers. He holds a Master's Degree in Computer Science from Syracuse University and an MBA from Duke University. His areas of expertise include designing e-business systems.

**Sungsim Park** is a Consulting IT Specialist and Solution Assurance manager in Technical Sales Support in IBM Korea. She has over 19 years of experience working on iSeries™, e-business applications and patterns, and Linux. She has taught several e-business on demand classes to IBM employees and Business Partners in Korea.

**Hugues Raguet** is an IT Architect in e-business solutions Technical Sales Support at IBM France. He has 9 years of experience in the IT industry. He holds a degree in electronic engineering from the ENSEEIHT University in Toulouse. His areas of expertise include the design of e-business solutions.

**Jack Schneider** is a Senior Consulting IT Architect in IBM US. He joined IBM in 1987 and has a consulting background with over 25 years of experience in architecting, designing, and implementing leading edge technology solutions to address customer's complex business needs.

**Laurent Vanel** is a System Architect in IBM France. He started as an IT specialist in the pSeries brand in 1990. His IBM experience included a 3-year assignment at the Austin ITSO. In his current position as a System Architect, Laurent is technically responsible for the infrastructure proposed to a selected set of large customers.

Thanks to the following individuals from across IBM who assisted with the development of the first edition of this redbook:

Alain Roessle	Alistair Rennie
Angel Luis Diaz	Annette Miller
Bala Rajaraman	Ben Amaba
Bruce Benfield	Bruce MacKenzie
Cathy Parker	Connie Nelin
Dan Wolfson	Elisabeth Poigin
Errol R Denger	Fabio L Marras
Frank Kyne	Fredrik Carlegren
Gerard Laumay	Georgina Castanon
Howie Miller	Jason Boxer
Jesus R Grana	Joe C Snyder
John Arwe	John Pershing
Joseph Labriola	Juergen Schneider
Julie Schuneman	Kathleen Harrigan
Kevin Larsen	Madeline Nick
Marc-Thomas Schmidt	Marjorie Macintyre

Mark Cathcart  
Matthew P Haynos  
Paola Bari  
Phil Fritz  
Randy B Daniel  
Robert Spory Jr  
Scott Vetter  
Thomas Lump  
Yann Guerin

Mark Palmer  
Miles Barel  
Pascal Durazzi  
Ralf Heindoefer  
Robert Liburdi  
Roy Bowen  
Simon K Johnston  
Tony Pearson  
IBM Learning Technologies at a  
Glance team

Our apologies to any others that may have been inadvertently left off this list.  
Your help was appreciated.

## Become a published author

Join us for a two- to six-week residency program! Help write an IBM Redbook dealing with specific products or solutions, while getting hands-on experience with leading-edge technologies. You'll team with IBM technical professionals, Business Partners and/or customers.

Your efforts will help increase product acceptance and customer satisfaction. As a bonus, you'll develop a network of contacts in IBM development labs, and increase your productivity and marketability.

Find out more about the residency program, browse the residency index, and apply online at:

[ibm.com/redbooks/residencies.html](http://ibm.com/redbooks/residencies.html)

## Comments welcome

Your comments are important to us!

We want our Redbooks to be as helpful as possible. Send us your comments about this or other Redbooks in one of the following ways:

- Use the online **Contact us** review redbook form found at:

[ibm.com/redbooks](http://ibm.com/redbooks)

- Send your comments in an Internet note to:

[redbook@us.ibm.com](mailto:redbook@us.ibm.com)

- Mail your comments to:

IBM Corporation, International Technical Support Organization  
Dept. JN9B Building 003 Internal Zip 2834  
11400 Burnet Road  
Austin, Texas 78758-3493





# Summary of changes

This section describes the major technical changes made in this edition of the book and in previous editions. This edition may also include minor corrections and editorial changes that are not identified.

Summary of Changes  
for SG24-6634-01  
for On Demand Operating Environment: Managing the Infrastructure  
as created or updated on June 2, 2005.

## June 2005, Second Edition

This revision reflects the addition, deletion, or modification of new and changed information described below.

### **New information**

This book now contains information about the IBM Virtualization Engine that was not announced when the previous version was published.

In particular, Chapter 2 and Chapter 3 describe what the IBM Virtualization Engine is. Part 2 of this book has been largely reviewed to include the IBM Virtualization Engine. Part 3 of this book has been extended with new examples in Chapter 17, Chapter 18, and Chapter 19.

### **Changed information**

Many changes and improvements were made throughout the book.





# Part 1

## Overview

In this part of the book, we introduce the On Demand Operating Environment framework and its infrastructure management component. These chapters introduce the reader to the components of an On Demand Operating Environment and the capabilities required to simplify its management.

- ▶ Chapter 1, “Introduction” on page 3 introduces the On Demand Business concept
- ▶ Chapter 2, “Infrastructure management overview” on page 13 introduces the infrastructure management component of the on demand concept
- ▶ Chapter 3, “The IBM Virtualization Engine” on page 45 describes one of the components of the on demand infrastructure capabilities, the IBM Virtualization Engine, that has been announced on August, 17, 2004





# Introduction

As enterprises strive to meet their current challenges, some of which are generally described in this section, they require an IT infrastructure that supports their business goals. This redbook and its companion (*On Demand Operating Environment: Creating Business Flexibility*, SG24-6633), describe an IT infrastructure that enables business to be more responsive, variable, focused, and resilient. Enterprises that exhibit these four attributes are moving to what IBM calls On Demand Business. The IT infrastructure described here and that supports an On Demand Business is called an On Demand Operating Environment.

In these redbooks we provide an overview of the components of an On Demand Operating Environment and we show how enterprises can start putting pieces in place today that address immediate business needs. Over the longer term, enterprises can build on these initial steps to evolve their IT infrastructure to an On Demand Operating Environment that enables the business to focus on their core competencies and meet whatever new challenges may arise.

# 1.1 Getting to On Demand Business

To ensure a successful IT project, it is critical to serve the right business priorities. Assessing business needs will lead down one of two paths: either customers will want to address their business design first, or their IT design.

Figure 1-1 describes the two approaches from IBM.

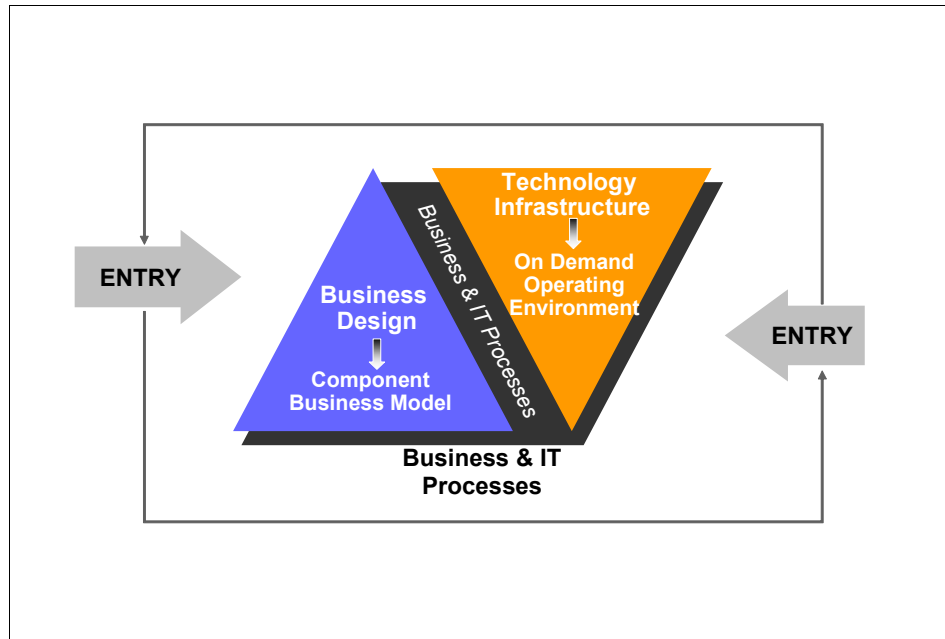


Figure 1-1 Positioning the On Demand Operating Environment

As shown in this diagram:

- ▶ For companies whose business goals tell them to start with business design, the Component Business Model™ will help identify and take aim at the right targets.
- ▶ Companies that already have a strong business model may need to focus on their technology infrastructure to optimize that business model. If you want to start with your technology infrastructure, IBM's On Demand Operating Environment concept applies.

On Demand is not about technology for the sake of technology, it's about enabling new ways of doing business. It's about helping an organization reach new levels of innovation while continuing to deliver the improvements in

productivity necessary to improve the bottom line. Yet the underlying technology makes an On Demand Business fundamentally different.

When business processes have been integrated end-to-end, across a company and with its key partners, suppliers, and customers, it has the ability to respond to any customer demand, market opportunity, or external threat. Yet, there's a lot of work to be done. Today's infrastructure is complex and rigid. Because much of it is based on proprietary hardware and software, delivered well before industry standards were established, it's difficult to make all the pieces work together. It's even more challenging to make them deliver the flexibility necessary to support today's dynamic business environment.

The need for change is forcing the emergence of a new computing model. This new on demand model blends the robust nature of the traditional IT computing model with the industry standards based computing model that enabled the Internet and the Web. It transcends both models, in a number of ways.

- ▶ The traditional IT model was focused on calculations, data processing, transactions, and other highly structured tasks. It served businesses well for those rigid applications and will continue to do so over time. But the model breaks down when trying to extend it into applications or processes that aren't so highly structured, such as long-term enterprise resource planning projects.
- ▶ The Internet computing model had a different design point. It gave us simple mechanisms, based on industry standards, to link together many components, which can be used to perform relatively simple functions such as browsing and searching for information, and sending and reading e-mail. The Internet computing model enabled a handful of new business models. But more important, it revolutionized the way that existing things were done, especially in the areas of communications between companies, marketing, sales, and customer support.

With that revolution came the recognition that computing technology is exponentially more powerful when it's based on industry standards. That meant the industry would need additional standards and mechanisms to handle more sophisticated applications.

The On Demand Operating Environment, as a computing model, builds on both the traditional and Internet computing models, leveraging industry standards to redefine how existing systems and technologies interact. This enables the creation of a highly modular environment, where application and infrastructure components can be more easily defined and managed. This allows for a more flexible and real-time implementation of business policies than was possible with more structured computing models.

We recognize that this isn't a one-size-fits-all solution or methodology. Organizations have different priorities, different personalities. An on demand

approach reflects that. With many different entry points, where to get started depends on the organization's priorities and resources.

In today's pragmatic environment, there are only a handful of organizations prepared to tackle all of the facets of creating an On Demand Business. Most companies opt to start more slowly. They focus on one key process and transform it. Or they start by taking steps to simplify their operating environment, increasing overall flexibility and resilience, while reducing the resources that their current approach requires.

## **1.2 Infrastructure to support an On Demand Business**

Over the last several years, most enterprises have concentrated on making individual business processes more efficient. This work has typically been done within application or line of business silos. Going forward, a continued improvement in business performance will require a horizontal view, looking across the business and even across the ecosystem of suppliers, partners, and customers.

To create applications and support business processes across lines of business or organizations will require the ability to use and integrate existing applications and processes. This will provide flexibility to allow the business to easily adapt and assemble new applications to support new business requirements. If there was ever an argument for using industry standards, this is it: enabling the business to quickly and seamlessly integrate processes that weren't built to work together, from a variety of vendors. With industry standards, applications don't need to be recreated every time some piece of hardware or software changes or is rewritten to support changes in the dependent processes.

Aside from the business flexibility that comes from the ability to integrate people, processes, and information across the business, the IT infrastructure must also be made simpler and more manageable. This includes support for virtualizing the resources required and automating the management and operations of the IT environment.

The characteristics that enable an on demand environment are the capabilities that enable business flexibility and simplification of the underlying technology infrastructure.

- ▶ The first focus is to increase business flexibility through capabilities designed to speed integration initiatives. The ability to connect people, processes and information in a way that allows the organization to become more flexible and responsive to the dynamics of its markets, customers and competitors is critical. It becomes increasingly so as the value net is extended to more tightly integrate partners, suppliers and customers into the business processes.



- ▶ The second focus is IT simplification, the creation of an infrastructure that's easier to provision, deploy, and manage. This is accomplished through the creation of a single, consolidated, logical view of, and access to, all available resources in a network. Many organizations have become comfortable with the practice of over-provisioning, buying excess capacity so they can handle the occasional spikes that almost every system experiences. Eliminating the practice of over-provisioning by moving to an infrastructure that accommodates dynamic resource provisioning can reduce an organization's capital investments significantly.

To achieve more flexibility and componentization, the infrastructure must evolve from silos of complex, over-provisioned, proprietary hardware and software to an industry-standards-based infrastructure, where capacity can be optimized across the entire organization.

## 1.3 Capabilities

On Demand Operating Environment capabilities enable business flexibility and IT simplification. As described in Figure 1-2, there are multiple key capabilities:

- ▶ Business-driven development across the development life cycle, to transform the traditional development into a business process. This requires rapid business prototyping, discovery of re-usable assets, development and test, deployment, all working together towards a common business goal.
- ▶ Integration, to connect people, processes and information. When everything connects in a Service Oriented Architecture, the business gains flexibility and insight.
- ▶ Infrastructure management, to simplify and to optimize operational IT services, through automation of provisioning, orchestration, and virtualization.

The overall objective is to evolve to an industry-standards-based, integrated, automated and virtualized IT environment.




Figure 1-2 The On Demand Operating Environment

Each of the capabilities of an On Demand Operating Environment acts as a facilitating element to enable the deployment of an underlying infrastructure that drives business flexibility and IT simplification. These capabilities are enabled by hardware, microcode, operating systems, middleware, and services. They need to be architected through a framework in a services manner and integrated seamlessly.

### 1.3.1 Integration capabilities

Integration capabilities enable the connection of people, processes, and information in a way that allows businesses to become more flexible to the dynamics of the markets, customers, and competitors around them. To maximize the ability to integrate within and beyond the enterprise, there are several key capabilities required. These will typically be implemented over time according to the needs of the individual business:

- ▶ *Business modeling* enables the graphical depiction and simulation of a business process, including task descriptions, resources required, and decision points.
- ▶ *Process transformation* enables existing applications and information to be reused in new ways.
- ▶ *Application and information integration* enables multiple information sources and business applications to be combined.

- ▶ *Access* extends data and information to new classes of devices and methods of interaction regardless of connection type.
- ▶ *Collaboration* allows users to interact in a personalized way with dynamic information, applications, processes, and people.
- ▶ *Business process management* can model, deploy, and analyze processes with the goal of managing the end-to-end business process.

### 1.3.2 Infrastructure management capabilities

Infrastructure management capabilities extend access to, and create a consolidated, logical view of resources across the network. This dramatically simplifies the operating environment, increasing flexibility and delivering broad-based cost savings. Fundamental to this simplification are the concepts of automation and virtualization.

To achieve this simplified and optimized management of the infrastructure, the following capabilities are required. Again, these will typically be achieved over time as the business requires:

- ▶ *Availability* helps ensure the health and appropriate functioning of IT environments.
- ▶ *Security* helps ensure that information assets, confidentiality, and data integrity are protected.
- ▶ *Optimization* helps make the most productive use of the IT infrastructure.
- ▶ *Provisioning* makes the right resources available to the right processes and people at the right time.
- ▶ *Policy-based orchestration* senses, triggers, and responds according to business goals.
- ▶ *Business service management* helps to visualize the IT environment in business terms and manage service levels to business objectives.
- ▶ *Resource virtualization* provides a single, consolidated, logical view of, and easy access to, all available resources in a network (including servers, storage, and distributed systems).

#### Automation

Automation enables an IT infrastructure to manage many day-to-day tasks itself. With a self-managing infrastructure, efficiency is increased and resource allocation simplified. A fully automated IT infrastructure can sense changing conditions, such as surges in demand or isolated application errors, and can spot trends that could lead to costly system downtime. The infrastructure automatically responds by taking corrective actions that ensure IT resources remain aligned with business goals.

## Virtualization

Virtualization is the ability to separate the direct dependency of an application to a physical resource. Through virtualization, an enterprise will:

- ▶ Have a single, consolidated view of, and easy access to, all available resources in the network, regardless of location.
- ▶ Efficiently access and manage those resources to reduce operations and systems management costs while maintaining needed capacity.
- ▶ Respond dynamically to the application needs of its users.
- ▶ Gather and access information across the organization quickly to gain competitive advantage.

Although we discuss the capabilities through the two entry points of integration and infrastructure management separately, in reality they are tightly linked. Security, for example, permeates IBM solutions, providing a critical, pervasive functionality across the On Demand Operating Environment.

## 1.4 On Demand Operating Environment architecture

An On Demand Operating Environment is an integrated infrastructure aligned to business goals and processes in a resilient and secure manner.

Figure 1-3 represents the On Demand Operating Environment architecture framework; the upper part describes the integration capabilities and the lower part describes the infrastructure management capabilities as defined in Section 1.3, “Capabilities” on page 7.

The *on demand computing model* applies at various levels in the IT stack.

- ▶ At the system level, the components are system objects (for example, computing capacity, storage, and files).
- ▶ At the application level, components are dynamically integrated application modules that constitute sophisticated, yet much more flexible applications.
- ▶ At the business level, the components are business objects, defined for particular vertical industries or more generally, as they apply horizontally across industries.

Because the on demand computing model is based on industry standards, it can be used to define the business, applications, and systems at various levels: within a department, across an entire enterprise, or throughout an industry ecosystem. It enables true end-to-end business process integration.

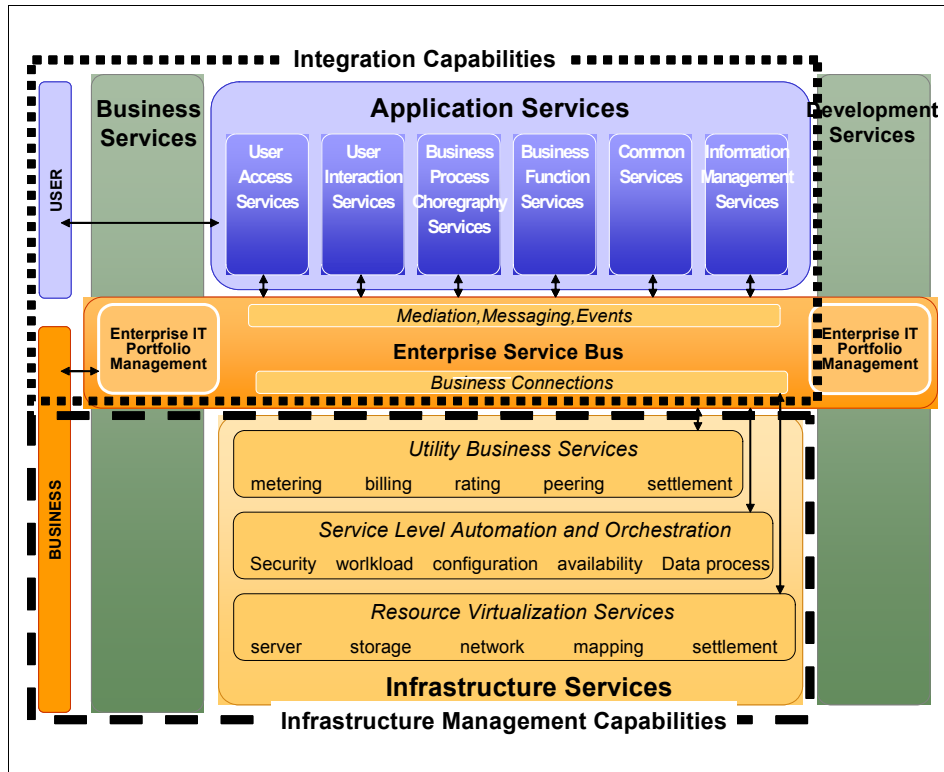


Figure 1-3 On Demand Operating Environment architecture

The *On Demand Operating Environment* is based upon the concepts of a Service Oriented Architecture (SOA). A Service Oriented Architecture views every application or resource as a service implementing a specific, identifiable set of (business) functions. Services communicate with each other by exchanging structured information—messages or documents (sometimes called business objects). Their capabilities are defined by interfaces declaring messages they can produce or consume, policy annotations declaring quality of service required or provided, and choreography annotations declaring behavioral constraints that must be respected in service interactions.

The actual implementation is hidden from the requester of a service, thus Service Oriented Architectures are a convenient way to achieve application integration by allowing new and existing applications to be quickly combined into new contexts. Existing applications are “adapted” to service declarations. The interaction of services can be direct, or can be mediated through an intelligent infrastructure, which we will call an Enterprise Service Bus (ESB).

Service Oriented Architectures require standards for the definition of services and their capabilities and interactions. The adoption of this architectural approach has been greatly facilitated by the growing acceptance of XML use to provide a standard representation of structured information and of “Web Services” standards (often called WS-\* standards). The conceptual model of a Service Oriented Architecture applies to the virtualization of both business functions and physical infrastructure. It spans the construction of applications as well as their deployment and management. A client (user or business) only sees a collection of business services, and is interested in their quality of service, but is shielded from the details of application assembly and service delivery through the On Demand Operating Environment.

## 1.5 Summary

In this chapter we have introduced the On Demand Operating Environment and briefly described its key capabilities, characteristics, and architecture. An On Demand Operating Environment is not a specific product or suite of products, and it is not something that will be created or deployed overnight. Enterprises will evolve to the Operating Environment by deploying various capabilities based on the specific needs of their business.

An On Demand Operating Environment provides businesses with flexibility by enabling integration between people, processes, and information, and creating a manageable infrastructure through automation and virtualization.

In the next chapter, we describe the infrastructure management capabilities in more detail.



## Infrastructure management overview

This chapter describes the infrastructure management capabilities of the On Demand Operating Environment framework. It is intended to provide the reader with an understanding of this portion of the framework as well as details on each of the sub-components within it. It describes the role each component plays within the framework, and the overall value and benefit each brings in positioning an IT organization's infrastructure to support an on demand strategy which addresses the business drivers the company faces.

In some cases, we include references to IBM products to help clarify these concepts. However, the specific product references provided should not be considered inclusive or prescriptive. As the On Demand Operating Environment will evolve, the list of functions, components, and products will continue to grow.

## 2.1 Business drivers

When evaluating how IT can support the business requirements of an enterprise, several questions are often asked by the CIO:

- ▶ How can I gain the flexibility to deploy new applications and systems faster in support of the organization's business needs?
- ▶ How can I more effectively manage the growing heterogeneous IT distributed application environments needed to support the business?
- ▶ How can I achieve higher return on investment through improved utilization of existing IT resources?
- ▶ How can I position the IT Infrastructure so that it can be responsive and flexible while maintaining its resilience and quality of service?

In many of today's organizations, business success is often tied to the speed with which business strategies can be changed to counter competitive pressures and capitalize on opportunities. Businesses require optimized integration of their people, processes, and information in order to have the ability to implement end-to-end business processes in support of business goals. This transformation into an On Demand Business requires building a dynamic infrastructure that can support these business processes. An infrastructure is required that is flexible and efficient, and that can be managed based on supporting the demands of the business. The IBM On Demand Operating Environment Infrastructure Management framework provides the components and technologies to help IT organizations position themselves to be able to manage the infrastructure required of an On Demand Business.

The business drivers imply a requirement to provide and operate a flexible and responsive IT infrastructure while controlling its associated costs. With the appropriate pieces in place this can be done through:

- ▶ Automation - freeing up human resources required to sustain the current environment so they can be used develop and deploy new systems driving new business
- ▶ Virtualization - allowing for optimized use of existing systems and associated resources, driving higher utilization and less incremental investment for new business applications

## 2.2 Framework for infrastructure management

From an infrastructure point of view, the customer maturity level, may vary, and solutions to be more effective can be developed depending on what is the primary business focus.



<b>CUSTOMER Type I</b>	<b>CUSTOMER Type II</b>	<b>CUSTOMER Type III</b>	<b>CUSTOMER Type IV</b>
<b>Primary focus on COST REDUCTION</b>	<b>Primary focus on EFFICIENCY</b>	<b>Primary focus on OPTIMIZATION</b>	<b>Primary focus on QoS and ROI</b>
Server to Customer Ratio is 1:1	Server to Customer Ratio is Many : 1	Server to Customer Ratio is Many : Couple	Service Model Based on Shared Resources
Islands of Computing	Islands of Computing	Negotiation between Departments	Service Offerings
Chargeback based on Per Server Pricing	Chargeback based on Per Image Pricing	Chargeback based on Usage and Cost of Nothing Index	Supply Chain Management
IT Focus on Recovery and Availability	IT Focus on Recovery, Availability and Performance	Policy Based Computing	Virtualized Resources
		IT Focus on IT Service Level	IT Focus on TCO/TCA and SLA

*Figure 2-1 Customer Infrastructure progression*

Depending on what issues customers are dealing with today and which problems they are facing in their environments, we can categorize the customers in four types, as described in Figure 2-1:

- ▶ Type 1 are focusing IT processes on recovery and availability.
- ▶ Type 2 are focusing on recovery, availability and performance.
- ▶ Type 3 are focusing on IT service level.
- ▶ Type 4 are focusing on the total cost of ownership (TCO).

Solutions to improve the maturity level will depend at which level a specific environment can be categorized. For example:

- ▶ To be more efficient, some physical consolidation may be investigated.
- ▶ To be more optimized, an enterprise-wide integrated strategy can be undertaken.
- ▶ To be more effective, the utilization of servers must be balanced according to business objectives.

In this redbook we look at infrastructure management from both an automation and a virtualization standpoint. These capabilities of the framework provide for an IT infrastructure that will help companies manage their heterogeneous environments, increase the utilization of the systems within those environments, provide for the flexibility to deploy and implement new IT resources dynamically as the business requires, manage the quality of service to SLAs, and reduce overall IT operating costs. The objective is to implement an IT environment that is efficient, effective, and responsive to the needs of the business.

The capabilities of the infrastructure management framework, as shown in Figure 2-2, can loosely be split into the following components.

- ▶ Automation components include:
  - Business Service Management
  - Policy-based Orchestration
  - Availability
  - Security
  - Optimization
  - Provisioning
- ▶ Virtualization components include:
  - Server Virtualization
  - Storage Virtualization
  - Distributed systems/grid computing
  - Network Virtualization

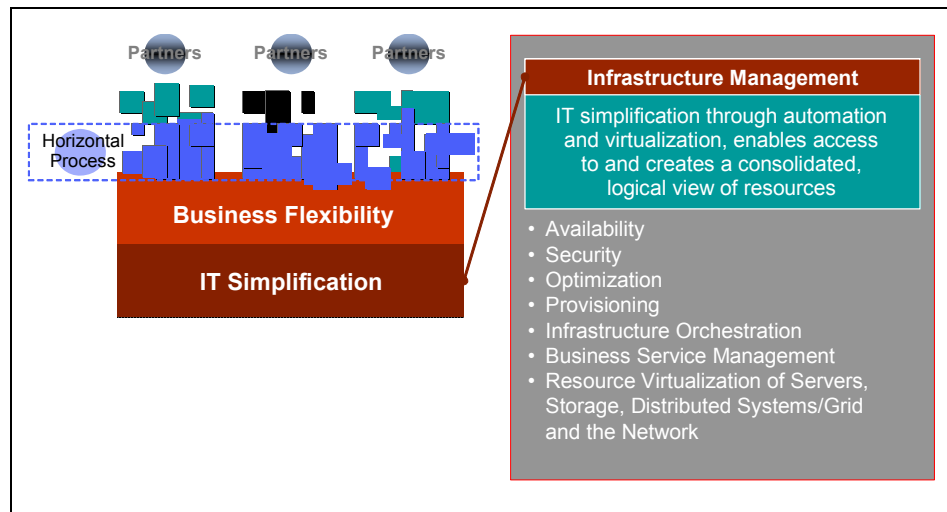


Figure 2-2 Capabilities for IT simplification

Each of these areas and their components are described in more detail in the following sections.

## 2.3 Automation

Infrastructure automation reduces the complexity of management, enables better use of assets, and improves availability and resiliency. Within the framework, automation consists of four elements: availability, security, optimization, and provisioning. These elements are linked together through Policy-Based Orchestration, which provides a high level of automation to coordinate IT resources to support the implementation of business policies across the infrastructure. In order to successfully automate the management of the environment to meet business goals, those service level objectives must be defined and available to the automation component of the framework. The Business Services Management component is where those goals are defined.

The highest level of automation is achieved through autonomic computing technologies. Autonomic computing provides the technology to enable information systems to be self-managing. These self-managing characteristics combine to deliver the automation required of an On Demand Operating Environment. The IBM Autonomic Computing Initiative focuses on developing autonomic technologies that enable the infrastructure to become self-managing. The core capabilities include:

- ▶ A solution knowledge architecture that addresses the change/deployment problem of any solution (application software, operating system) in a non-homogeneous environment. By capturing installation and configuration information in a consistent manner, a common solution knowledge capability eliminates the complexity introduced by many formats and many install tools.
- ▶ A common system administration, whose primary goal is to provide a single platform that can host all the administrative console functions for the server, software, and storage. A common console instance consists of a framework and a set of console-specific components provided by other product development groups.
- ▶ A problem determination architecture that normalizes the data collected, in terms of format, content, organization, and sufficiency (known as the Common Base Event format). The architecture defines an adapter/agent infrastructure that provides the ability to plug in adapters to transform data from logs and traces to the standard format, as well as sensors to control data collection (filtering, aggregation, correlation, and so on).
- ▶ An autonomic monitoring capability that provides a runtime environment for a resource manager to gather and filter data obtained through sensors. This capability includes a common way to capture the information, a set of pre-defined resource models, a way to incorporate policy knowledge, and a way to plug in multiple engines (for event isolation, for basic root cause analysis and server-level correlation across multiple IT systems, and to automate initiation of corrective actions, for example).

- ▶ A rule language that supports reasoning through procedural and declarative rule-based processing of managed resource data used for complex analysis. The underlying rule engines can be used to analyze data. Application classes can be imported directly into rule sets so that data can be accessed (using sensors) and control actions can be invoked directly from rules (using effectors).
- ▶ A uniform method for defining the policies that govern the decision-making. By defining policies in a standard way, they can be shared across resource managers to enable entire systems to be managed by a common set of policies.
- ▶ A clear understanding of the transaction topology that maps the service classes to this topology. Autonomic managers involved in workload management need a transaction measurements capability to understand how the systems involved commit their resources to execute the workload, and how changes in allocation affect performance over time.

Working in a heterogeneous distributed environment requires the deployment of autonomic managers throughout the IT infrastructure, implying a diverse range of suppliers. These systems, therefore, must be based on open industry standards. New open standards are being developed and shared with the industry that will define the new mechanisms for interoperating in a heterogeneous system environment. This is one of the main missions of the IBM Autonomic Computing Initiative.

Autonomic computing is a fundamental underpinning for automation offerings.

### **2.3.1 Business Service Management**

IT organizations today focus their attention on application availability and performance metrics. As a company evolves into an On Demand Business, the focus will move to determining how the business and business processes are impacted from the infrastructure resources that have been applied and expended to service the specific workload. These services allow IT organizations to define service level goals as they apply from a business standpoint. The services then track, report, and price the on demand resources that are utilized in meeting those goals.

In the past, IT organizations have always struggled with capturing costs and services that were associated with the delivery of resources for a specific business function. Service level objectives for the business application running within the framework are established and monitored based on the business goals and requirements defined at this level of the framework. Business Service Management monitors, meters, tracks, and provides capabilities to handle accounting for charging back infrastructure costs to business units.

## 2.3.2 Policy-based Orchestration

Policy-based Orchestration plays a major role within the infrastructure management framework by providing policy-based Quality of Service (QoS) delivery coordination across the four core automation disciplines of the On Demand Operating Environment (availability, security, optimization, and provisioning). Resource allocation can be performed to enable the creation of heterogeneous n-tiered application environments from pooled resources. In most cases, implementations of business applications involve many different tiers and technologies on each of those tiers. The orchestrator has the responsibility of monitoring the environment, analyzing its data, and then determining from established policies what actions or workflows need to be performed in support of satisfying SLA objectives that were established for the business application.

### **Enforcement of business policies for automated change**

The orchestration component understands and maintains application service level objectives as they relate to response times and availability. It contains the knowledge base to make the decisions as to what types of resources are to be provisioned in support of a specific business need based upon the SLAs that have been established. For example, if it is determined that there is too much network traffic for the existing set of Web servers, a new Web server needs to be provisioned. The orchestration component knows what platform and operating system should be used to provide the required level of service. It understands the Service Level Agreements and IT policies for deploying new Web servers and will know whether the new resource should get provisioned on a blade center, for example, and what operating system should be used.

### **Automated allocation of resources to prioritized applications**

Meeting SLAs to support the needs of the business is critical within an On Demand Operating Environment. To ensure that the most important application's service levels are achieved during periods of resource contention, multiple "classes of service" are supported. As an example, just as the orchestrator has the ability to direct the provisioning of new resources in an effort to maintain service levels, it also has the ability to de-provision resources. In certain situations, the orchestrator component may need to make decisions regarding when resources need to be de-provisioned from lower priority applications so that the resources can then be re-provisioned to meet the needs of higher priority business applications.

For example, on zSeries, Workload Management (WLM) can manage the performance and number of application servant regions in WebSphere Application Server for z/OS. WLM manages the response time and throughput of WebSphere transactions according to the goals assigned in the WLM service policy and the availability of system resources. To meet these goals, WLM

controls the number of active servant regions, ensuring that sufficient processing capacity is available to deliver the required service levels.

These same concepts will also apply to distributed systems, allowing workload to be dynamically allocated to available systems, and systems dynamically allocated to applications.

### 2.3.3 Availability

Capabilities in the availability component are focused on the resiliency of the environment. This addresses areas related to disaster recovery mechanisms, resource provisioning, data replication, transaction management, event management, and root-cause analysis.

#### Transaction management

Having a server provisioned, operational, and available does not necessarily imply that the application on that system is available or that a transaction can be completed. In an On Demand Operating Environment, where multi-tiered implementations can span various heterogeneous servers, and storage and network resources, a more holistic view of applications must be considered to ensure that an end-to-end transaction can be executed successfully in support of the business. To achieve this state, the availability and orchestration components need to work together to ensure that all the pieces that make up an application provide the service. Modelling end-to-end applications as a logical composition of resources, and by defining operational dependencies between those, will help autonomic managers to also automate the recovery or failover process needed in case of planned and unplanned outages.

An example of the synergy between Availability and Orchestration is in the zSeries environment, where Parallel Sysplex in tandem with the WLM service policy, provides the ability for applications to be spread across a number of operating system instances, protecting the application from the outage of a system. Geographically Dispersed Parallel Sysplex™ (GDPS®) builds on this capability, providing the potential to protect application and data availability across planned and unplanned system, DASD, and even complete site outages. Currently, GDPS can manage the remote-copied Open LUN<sup>1</sup> (Logical Unit Number) volumes used by non-z/OS systems; in the future, agents will be provided for those systems to allow them to initiate GDPS operations should they detect a failure.

---

<sup>1</sup> The Logical Unit Number (LUN) is introduced in Section , “From a point-to-point topology to a network topology” on page 35

## **Event management and root-cause analysis**

The objective in an On Demand Operating Environment is to have autonomic systems which are self healing. Within on demand infrastructures with large, multi-tiered, heterogeneous business application implementations, identifying and isolating the cause of a problem can become very complex. Event management and the automated correlation of events occurring across the tiers of the infrastructure become necessary in order to quickly address and resolve problems that are impacting systems availability or have the potential to impact that availability. Based on this fault isolation and root-cause analysis steps, automated recovery action can be concluded based on the end-to-end application topology model. This is particularly true when the application topology defines some high availability redundancy which could be used to provide the broken service again.

## **Dynamic sense and respond**

The infrastructure must have the ability to sense what is happening across the entire environment and know its impact to the applications (and their impact on the business) in order to make decisions about how to respond. For example, let's have a look at an Internet banking application. If a Web server goes down, based on business priorities, a certain level of importance gets associated with the situation and a determination can be made about how to respond. In this situation, the optimum solution might be to provision a new server to handle the transactions without the end-user or consumer knowing about the change. In the case of a critical database crash, a high availability product is needed for a policy-based automated failover of the database to a backup node with minimal downtime. In both cases, the end-user is fully insulated from the underlying infrastructure that is handling their banking transactions.

Based on the recovery time objective and the business importance, a combination of server provisioning and traditional high availability products such as HACMP™ or the Tivoli Systems Automation products need to cooperate as well new technologies such as Enterprise Workload Management and Cisco solutions. With this combination it is possible to allow a fast recovery using the traditional high availability products, with the augmentation of dynamic server deployments to support availability objectives automatically.

### **2.3.4 Security**

Every IT organization is faced with managing the security of many different resources within their enterprise. The security component of the framework focuses on the distributed security mechanisms needed to support and manage these resources. Mechanisms are needed for authentication, authorization, user lifecycle management, privacy management, and directory services, within and across organizational boundaries. To effectively secure an On Demand

Operating Environment where resources are automatically and dynamically provisioned, based on the demands of the business, IT organizations need an end-to-end security infrastructure that can support such an implementation. Evolving inter-operability and industry standards will be key to allowing effective security implementations within an organization's environment as well as outside of that environment. These implementations will support integrated identity management, including: access control management, identity management, and privacy management.

### **Identity management**

On Demand Business needs to get to information quickly and easily in order to compete in today's marketplace. The ability to manage access to the growing number of resources within IT organizations and provide that access to an ever-expanding community of users, is very challenging. Moving to an on demand environment where resources are dynamically provisioned requires timely and accurate access to those resources for the right people within the right organization. Integrated identity management addresses the need for this ability with a centralized approach to managing the policies related to controlling access to resources.

### **Access control**

Applications and data must be secured dynamically. Access control is needed to provide consistent policy-based access to data and applications across the enterprise. The service must support multiple authentication mechanisms from passwords to certificates to tokens.

### **Privacy management**

In addition to centrally coordinated identity management and access control, privacy management must also be centralized and must address the need to implement policies that govern sensitive information. In an On Demand Operating Environment, automation in applying these policies ensures consistent corporate compliance across the entire enterprise for all assets within the IT infrastructure.

## **2.3.5 Optimization**

The optimization component provides capabilities targeted at achieving the performance goals required to satisfy the objectives set to support the business. This includes functions such as capacity analysis, planning services, job scheduling, and workload management.

A couple of examples of optimization are described in the following paragraphs.



## Organization-wide storage usage

Optimization is not only focused on servers and applications, but needs to apply to all pieces of the infrastructure to ensure the organization maximizes the use of all resources and minimizes their costs. For example, a production application, based on projections, may have a certain amount of storage resources allocated to it. But for some period of time it will not actually be using the storage. If the storage allocation is fixed, other applications which may have an immediate need for storage will either suffer without the additional allocation or the company will need to procure additional storage to meet the needs of these other applications. If the storage for the production application was optimized, only the required storage would have been allocated. The remaining “free space” would then be available for other applications to use, thereby allowing the company to defer storage costs while still meeting the SLAs associated with the production application.

IBM has long experience with these concepts of storage optimization and application storage pooling — the DFSMS component and DFSMSHsm™ product that deliver this capability have long been available in the zSeries environment. Multiple aspects of the storage optimization infrastructure are described in this book.

## Web site intelligence

The optimization component also looks at other areas of the infrastructure for information that, when combined with the knowledge within the orchestrator regarding business service levels, can help in making determinations as to whether additional resources are needed in support of that specific business area. For example, if traffic to a certain Web site that is part of the infrastructure that supports a critical business function is increasing rapidly, the optimization component combined with the orchestrator would address that situation by allocating additional Web servers to allow the load to be balanced and response times to remain within established SLAs.

For example, WLM on the zSeries platform, working together with the subsystems and Communications Server, has been able to balance the different workloads and provide resources to applications as needed. WLM also works with WebSphere Application Server (WAS) on z/OS to start more servers if the workload warrants it. Enterprise Workload Management (EWLM), a component of the IBM Virtualization Engine, will be able to extend this capability to heterogeneous platforms, both IBM and non-IBM, and be able to manage resources on these servers. Already today EWLM has the ability to influence the workloads across multiple Web servers together with load balancers such as Cisco and Nortel (Chapter 17, “How to balance workloads in the network” on page 255 provides more detail for these functions).

## 2.3.6 Provisioning

Provisioning handles the coordination and allocation of managed resources such as servers or storage within an On Demand Operating Environment.

These resources are provisioned dynamically since they are required for satisfying the needs of a business application. Provisioning is done by allocating managed resources from a “resource pool.” This gives the IT organization the flexibility to dynamically respond to the needs of the business. Resources can be dynamically allocated from the pool, then de-allocated and returned to the pool after the resource is no longer needed for a specific business application. Once de-allocated and returned to the pool, the resource then becomes readily available to be allocated to another application where resources might be needed to ensure performance goals and commitments of the application are met.

In addition to provisioning hardware resources, operating systems images, middleware such as Web applications, server software, and DB software can be loaded by the provisioning component of the framework. The provisioning component provides the “what to do” piece of automation to ensure that a newly provisioned resource is made fully usable to support the business goals.

Provisioning and workload management work together to provide resources where needed. Workload management is responsible for managing the existing resources, such as the system and network resources that are available to meet (where possible) SLA goals.

Given a pool of servers, the provisioning component will attempt to maximize the effectiveness of the pool by increasing or decreasing the number of servers assigned to a given cluster.

For example, in the zSeries environment, Intelligent Resource Director works with WLM and LPAR to move CPU capacity and channel bandwidth to the systems running the most important workloads (as determined by the WLM service policy) at any given time. Capacity Upgrade on Demand and Daily On/Off Capacity on Demand provide the ability to non-disruptively add and remove additional CPU capacity, and Variable Workload Charges provide the ability to only pay software charges for the CPU capacity actually used, rather than for all the capacity that is installed.

### **Complete cross-resource environment**

The provisioning component will not only handle servers and storage but network devices as well. Provisioning and management of networking equipment such as routers and load balancers could also be required as part of an on demand infrastructure to address a scaling or performance issue. These devices might be

provisioned to improve network routing or to handle an increased load being experienced by a Web server or by a business application itself. These devices can also be returned to the pool when no longer required, and reused to meet service level requirements of other systems within the environment.

### **Application provisioning**

Provisioning can also be used to address the need to handle additional application workloads. Just as it is possible to add additional hardware devices such as Web servers to address peak demand, it is also possible to dynamically add, for example, additional Web application servers into an existing cluster to absorb an increased number of transactions related to the implementation of a new business process. Again, this ability is needed in an On Demand Operating Environment to ensure the needs of the business are satisfied.

It is currently available on the zSeries platform through the cooperation of WebSphere Application Server and WLM.

### **Identity mapping**

Once a resource has been provisioned, it needs to be secured and access to it needs to be granted to the appropriate users. The provisioning component of the framework handles identity mapping for the resource and works with the security mechanism within the infrastructure to ensure the appropriate access is provided.

## **2.4 Virtualization**

The automation concepts described previously, are more effective when some virtualization capabilities have been implemented. The virtualization architecture builds the base layer for the On Demand Operating Environment architecture. Automation and virtualization are both part of IT simplification and define the needed infrastructure management to start an On Demand Operating Environment.

Virtualization can be broken into four areas:

- ▶ Server virtualization
- ▶ Storage virtualization
- ▶ Network virtualization
- ▶ Distributed and grid systems

These sections are detailed below in the next chapters. Although virtualization is not a new technology strategy, it is receiving a lot of attention today.

## 2.4.1 The value of virtualization

There are many forms of mature hardware, software, and solutions in virtualization technologies that are available for exploitation today. Each of them contributes to the business values of cost of IT, quality of service, and time to value. In the following sections, we describe some of the pain points that the virtualization technologies want to address.

### **Fragmented heterogeneous infrastructure**

Today, most applications are not typically designed to be integrated with other applications or software. Applications are generally developed to support one specific business area or function. They have often been monolithic in design and, although there are exceptions, each application usually runs on its own dedicated physical server or set of physical servers, and has its own physical storage devices. There is very little, if any, sharing of IT resources between different applications unless they were deployed on mainframe-based machines or on machines that support partitioning. For the most part, non-mainframe applications get deployed on their own specific set of physical hardware and network resources. This has led to the fragmented heterogeneous infrastructures that exist today, and it often results in a lack of flexibility within an organization's IT environment.

### **Low utilization rate**

IT organizations struggle with the fact that the servers on which these applications are deployed usually have low utilization rates, and the unused system resources sit idle and are wasted. Costs for today's implementations are driven even higher as IT departments over-provision systems to ensure that processing power or storage is available to meet peak demands for the specific business area being supported. The associated technical and operational costs related to this type of environment continue to rise as the number of physical resources and complexity continues to grow.

### **The existing technologies**

Various and many technologies can be utilized already today to help businesses better leverage their servers, networks, and storage infrastructures.

For example, IT departments can partition some servers to increase their utilization using logical partitioning capabilities (as they exist today for the zSeries, the pSeries and the iSeries), or they can use Blade technology to enable better management of Intel-based systems, or they can take advantage of networking capabilities that are built into servers today, such as HyperSockets™ and Virtual LANs, or they can exploit the capabilities of current storage technologies such as block virtualization, file aggregation, and centralized management.

## The On Demand Operating Environment challenge

As businesses move forward with on demand strategies that drive implementations of true end-to-end integrated business processes, the number of application systems that support those processes will continue to grow and will increasingly span multiple servers on different technologies. These heterogeneous cross-platform technologies will need to be managed, monitored, and measured holistically to ensure that the needs of the business and SLAs are being met. In an On Demand Operating Environment, virtualization is needed to uncouple the applications from the physical configurations.

### 2.4.2 Server virtualization

A virtual server is an environment which seems to be real for the end user. It can be only one part of a high end machine partitioning) or it can be an assembly of several machines, doing the same workload (clustering). The following sections describe these two virtualization capabilities; the last section describes some of the server features which enhance these capabilities.

#### Partitioning

IBM has been delivering systems that have provided virtualization for many years. Workloads running on mainframes today exploit technologies such as *virtual memory*: Each application thinks it has its own real, dedicated memory, though in fact they are sharing the same memory securely and reliably.

*Logical partitioning* (LPAR) allows customers to “slice up” a machine into virtual partitions, and provides the flexibility to dynamically change the allocation of system resources for those partitions. Any of the virtual servers may run on any of the physical engines, meaning that the engine resources are fully shared. Using intelligent dispatching within the logical partitions, the shared resources belonging to the physical server can be more efficiently used, thus allowing very high utilization levels.

Three levels of partitioning exist: physical partitioning, logical partitioning and software partitioning, as described in Figure 2-3 and detailed in the following sections.

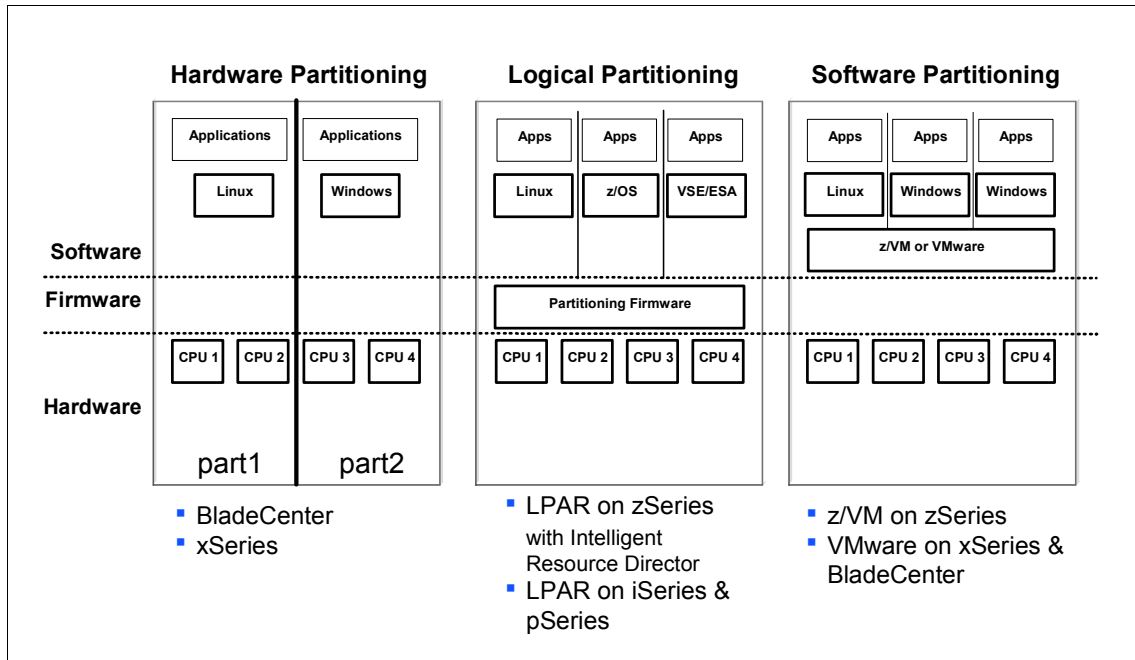


Figure 2-3 Types of server partitioning

### Physical partitioning

In physical partitioning the partitions are divided along hardware boundaries; the number of partitions relies on the hardware. Each partition might run a different version of operating systems. Each partition acts as a physically independent, self-contained server with its own processors, memory, input/ output subsystem and network resources. Physical partitions allow complete isolation of operations from operations running on other processors, thus ensuring their availability.

As shown in the “hardware partitioning” side in the Figure 2-3, the hardware box has been physically partitioned into two virtual boxes; the CPUs, the memory and the disks have been split between two partitions “part1” and “part2”. Part1 and part2 can be seen as two independent boxes with no access to each others resources.

Physical partitioning has been introduced in the mainframe more than 30 years ago and is part of the base mainframe architecture. Some Windows® and UNIX® servers support physical partitioning, in which each physical partition or virtual server owns a dedicated set of resources not used by others, and where the capacity granularity is based on multiple engines (2 to 4), because the virtualization is aligned with the physical components.

Virtualization granularity, at that level, does not generally provide significant advantages with respect to consolidation unless it is combined with other software types of virtualization methods. This is because each partition acts like a separate server, with dedicated resources for individual workloads. In this configuration, each server needs to have dedicated resources to accommodate peak utilization, hence low resource utilization for the majority of the time. The affinity of the virtual servers to the physical board configuration also means that resources such as processor engines, memory, and I/O interfaces have to be provisioned and re-provisioned together, thus making the implementation less efficient.

### ***Logical partitioning***

With IBM @server® Logical Partitioning, there is no affinity between the processor engine resources, the memory, and the I/O interfaces, thereby making it possible to provision and re-provision the resources to the virtual servers independently. As drawn in the “logical partition” side of Figure 2-3 on page 28, the logical partitioning implies an architecture implementation in the firmware. The CPUs, the memory, and the I/O resources can belong to any logical partition; these resources can be dynamically moved among partitions without operating system and application disruption. The hardware-based virtualization capabilities of the IBM mainframe also allow virtualization and sharing of I/O paths, and allows for the definition of virtual TCP/IP networks connecting the virtual servers at memory speed.

For example:

- ▶ By utilizing common hardware and microcode componentry, both the IBM @server® i5 and p5 systems offer Micro-Partitioning™ where the POWER5™ processors on the servers can be allocated to tenths of a processor among the partitions. In addition, both the iSeries and pSeries allow you to run multiple operating systems in different partitions, and allow processors, memory, and I/O to be shifted among active partitions without requiring the operating system to be rebooted.
- ▶ Logical Partitioning (LPAR) on zSeries servers is enabled by the Processor Resource System Manager (PR/SM™), a standard feature of all S/390® and zSeries servers. Introduced more than 20 years ago, this capability was designed to help in isolating workloads in different z/OS images, so we can run production work separately from test work, or even consolidate multiple servers into a single processor. An LPAR consists of a set of physical resources (CPU, storage, and channels) controlled by just one independent image of an operating system, such as z/OS, OS/390, Linux, VM (see “Software partitioning” on page 30), or VSE. Different parameters allow you to customize the environment: for example, the number of logical processors per LPAR, the LPAR weight (LPAR weights determine the minimum guaranteed amount of physical CP resource an LP should receive), the storage size

partition, and the LPAR capping (LPAR capping is a function used to ensure that a partition use of the physical processors cannot exceed the amount specified). You can have up to 30 LPARs in a zSeries complex today; in the future, 60 partitions will be possible. LPARs have achieved common criteria security certification Evaluation Assurance Level (EAL) 5 on the zSeries 990, 900 and 800 servers.

LPAR capping is a function used to ensure that an LP use of the physical CPs cannot exceed the amount specified in its Target (LPx).

### ***Software partitioning***

Software partitioning, as shown in the “software partitioning” part of the Figure 2-3 requires a specific software to be installed. On top of this software, multiple and different operating systems may be installed. The type and the number of operating systems that can be installed depend on the partitioning software itself. The number of partitions can be several hundreds. As with logical partitioning, the resources may not be allocated to a specific partition, and they can be moved dynamically between the partitions (this is an installation choice to decide if resources are dedicated, or not, to specific partitions).

For example:

- ▶ The most advanced software virtualization product on the market is the IBM z/VM® product running on the IBM zSeries. z/VM takes advantage of the hardware capabilities within the mainframe hardware and establishes a very low overhead and highly secure virtualization platform. It allows for the consolidation of large numbers of servers in a virtual server farm. z/VM is also the main virtualization enabler for Linux on the mainframe. Using Linux as a guest of z/VM allows you to run tens to hundreds of Linux images on a single zSeries server. z/VM V5 offers an ideal platform for consolidating select UNIX, Microsoft® Windows, and Linux workloads on a single physical zSeries server and for hosting other zSeries operating systems, including Linux for zSeries, z/OS, z/OS.e, VSE/ESA™, and z/TPF, as well as z/VM itself as guests.
- ▶ On xSeries® machines, VMware<sup>2</sup>, a non-IBM product, acts as a software hypervisor, hosted in a Linux kernel, and offers the management and the allocation of the dedicated or shared hardware components.

With these technologies, the installed operating systems in the “guest” images become independent with the hardware level. For example, no more specific drivers are needed. It is very easy to migrate, change, or consolidate any virtual machines you have.

---

<sup>2</sup> For more information, look at <http://www.vmware.com/>



## **Summary**

Hardware partitioning uses fixed resources; software and logical partitioning offer greater, more dynamic capabilities, depending on implementation and platform.

## **Clustering**

Clustering of servers provides a way to make multiple server resources, and in some situations multiple data resources, appear as a single resource space. The value of this is the ability to scale capacity and establish highly available solutions. Clustering can be used to “virtualize” the server or provide some of the cost benefits related to the server virtualization technologies previously described.

Some examples of clustering are the RS/6000® node clusters, the zSeries Parallel Sysplex, the Linux Cluster, the IBM Blade Center technology, the HAMCP implementation or the IBM Cluster System Management.

We notice that there are several ways to create a cluster, and the type of implementation can be of various kinds: hardware, built-in software or external management tools.

### ***With the Parallel Sysplex***

Parallel Sysplex clustering was designed to bring the power of parallel processing to business-critical applications running on a zSeries infrastructure. A Parallel Sysplex cluster consists of up to 32 z/OS images, coupled to one or more Coupling Facilities (CF) using high-speed specialized links for communication. The Coupling Facilities enable high-speed, record-level read/write data sharing among the images in a cluster. The CF can be standalone servers or a logical partition on the zSeries server. The CF allows high performance, multisystem data sharing across all the systems. In addition, workloads can be dynamically balanced across systems with the help of workload management functions (for more information about WLM for z/OS, please refer to *OS/390 Workload Manager Implementation and Exploitation*, SG24-5326). With the Parallel Sysplex, applications that support data sharing can potentially run on any system in the sysplex, thus allowing to move the workload to where the processing resources are available.

### ***With blade solutions***

Blade solutions provide for consolidation with virtualization benefits, where the physical space used is reduced and a limited amount of hardware infrastructure sharing takes place. Virtualization of individual blades is also possible. The processor resource space being shared is relatively limited, and this reduces the ability to move capacity dynamically between the different virtual servers. The IBM @server® product line provides efficient blade solutions both for xSeries supporting Windows and Linux and for pSeries supporting AIX and Linux.

The IBM BladeCenter™ can be used to increase the cost efficiency of these Intel-based servers by providing the capability to share system resources such as Ethernet adapters and fiber optic switches across all blades within the BladeCenter. It offers the capability to scale out. A BladeCenter can be purchased without all blade slots being fully populated. As additional capacity is needed, more blades can be purchased to populate the BladeCenter. This variable cost model allows the unit to be partially populated with blades, thereby reducing the number of blades that may sit idle in the BladeCenter until some future point when the business need drives application demand for the resource.

### ***With grid technology***

For many organizations, the combination of clustering, server virtualization, blade-based hardware and a services-based infrastructure will provide a compelling commodity infrastructure. Increasingly, clustering based on the successor to the grid standard, the OASIS Web Services-Resource Framework, and related industry standards, will be an increasingly attractive option as IBM delivers its On Demand Operating Environment.

Combining grid-based technology and heterogeneous SMP and/or blade-type servers that are co-located and interconnected by high-speed clustering — either vertically between like systems, or horizontally between like and unlike systems — will become an increasingly attractive way to run a heterogeneous utility-like infrastructure.

### ***Clustering value***

With clustering, you can use more of your computing power while ensuring that critical applications continue operations when recovering from a hardware or software failure. High availability clustering minimizes, or ideally, eliminates, the need to take resources out of service during maintenance and re configuration activities (for example, HACMP AIX software provides a fast recovery feature to minimize unplanned downtime). Another benefit is the ability to add and subtract servers from the cluster, which allows the cluster to scale to meet end-user demands for the service.

### ***Other features***

The following features are not directly part of partitioning or clustering, but from the hardware point of view, they can bring an added value to be implemented in a virtualized environment.

For example a virtualization capability, *Capacity on Demand*, is implemented across the IBM @server® platforms. zSeries, iSeries, and pSeries servers provide support for temporary or permanent processor enablement based on a client need. This feature does increase flexibility of the business side by reducing costs.

Virtualization is now expanded beyond the scope of one physical box and one architecture to deliver the classical benefits of virtualization on an end-to-end scale, across all architectures including servers, storage, and networks. This new level of virtualization is realized through technologies such as provisioning and workload management.

*Workload management* provides the control to ensure system resources are provided to the applications that are most critical to the business. Intelligent Resource Director (IRD) in zSeries machines expands on the workload management concept to provide goal-based resource balancing across multiple LPARs. Workload Manager on zSeries (WLM) and Partition Load Manager (PLM) in pSeries have the same kind of functions which balance hardware resources between partitions.

## Conclusions

IT organizations that exploit these technologies derive value and benefits in the form of easier management and cost savings from sharing system resources, which in turn results in more efficient use of those resources. Businesses can take action today to become more responsive by exploiting these same types of technologies as they continue to evolve into other IBM *@server®* platforms.

Note that increasingly there appear to be many similar technologies offered by a number of vendors of server virtualization technology. However, organizations need to take care in selecting an appropriate technology, because while they may provide “virtual machine” or partitions, many alternative implementations do not provide the same granularity, flexibility and isolation, although they may be marketed using the same terminology.

These are key values of any server virtualization technology and the more granular the technology is, the more flexible it is likely to be. Inherently the more flexible it is, the higher level of isolation it needs to deliver in order to achieve the reliability required to deliver on demand.

### 2.4.3 Storage virtualization

Data is the heart of any IT environment, all computing evolutions aim to increase the capability to use more and more data to achieve a specific target. All applications need access to data, handling various requests and generating responses and new data. The storage of data is therefore very critical.

At a business level, clients are faced with three major storage challenges with their storage infrastructure.

► **Managing storage growth:**

According to IBM studies, storage needs will continue to grow at over 50% per year in the coming years. Managing storage is becoming more complex than ever. We now have to deal with multiple server platforms and different operating systems, which may be connected to a Storage Area Network (SAN) with multiple and diverse storage platforms.

► **Increasing complexity:**

Although the declining cost of storage per megabyte (MB) makes it attractive to add additional disks, the increasing complexity of managing this storage results in over-used staff and under-used information technology (IT) resources. Combining this with the shortage of skilled storage administrators, it is possible to add significant cost and introduce risk to storage management.

► **Maintaining availability:**

The added complexity of 24x7 environments significantly reduces the efficiency of conducting routine maintenance, scheduling backups, data migration, and introducing new software and hardware. This problem is compounded by the fact that, as availability increases, so does the inherent costs.

Figure 2-4 positions the different architecture components of the storage infrastructure. In this chapter we focus on Storage Infrastructure Management which is the TotalStorage Productivity Center and Storage Virtualization which consists of the SAN Volume Controller and SAN File System. The following paragraphs introduce how the storage infrastructure evolved from a very direct server connection to a virtualization connectivity layer.




Figure 2-4 Storage Virtualization components

## From a point-to-point topology to a network topology

A Direct Attached Storage (DAS) is a storage device that connects directly to a single server. The devices are directly attached to a host computer through various adapters with standardized data transport protocols such as IDE (Integrated Drive Electronics) or SCSI (Small Computer System Interface). Network workstations may access server storage across a LAN network.

The most popular is the SCSI, which defines both a data communication protocol and a parallel communication bus for interconnecting devices with host processors. Though well performing, there are limitations in distance and sharing among users. Furthermore the server vendors implemented variations that are not fully compatible and there are some variations in the SCSI implementations.

For better optimization of the storage space and to address the limitations of the DAS architecture, a first level of physical consolidation was introduced with the Storage Subsystem or controllers such as the IBM TotalStorage Enterprise Storage Server® (ESS), the DS4000 family (FASTT) or the new IBM TotalStorage DS6000 and DS8000 series.

With this physical consolidation, a first logical level of virtualization was already introduced into the storage subsystems: the Logical Unit. Logical Units, which are the unit of storage that can be defined in these storage subsystems and assigned to the different servers, are in fact commonly called LUNs. Historically, a Logical Unit Number (LUN) is a unique identifier used on a SCSI bus that enables it to differentiate between multiple separate devices (each of which is a Logical Unit).

Because of the growing number of users and disk space to manage, the storage network was then introduced to offer a better accessibility, allow boundless distances and allow better devices connectivity and sharing; storage devices are no longer directly attached to the servers but accessed through a network topology.

## Storage Area Network and Network Attached Storage

The network capabilities addressed two kinds of data access:

- The *file level access* with the Network Attached Storage (NAS):

NAS is platform and operating system independent. The NAS uses a network topology such as a Local Area Network (LAN), supports multiple network file access protocols and is easy for maintenance. Many protocols can use it and present this data as if they were on the machine. NAS is optimized to move the data to users efficiently without the overhead and complexity of network servers.

- The *block level access* (or disk access) with the Storage Area Network (SAN):  
SAN connectivity is accomplished independently of the Local Area Network (LAN), using a high-speed network and protocol such as Fibre Channel or iSCSI (internet SCSI). SAN storage is accessible from the servers, so users can access any storage device on the SAN, regardless of the physical location of the storage or users. The SAN uses a special-purpose network separate from the LAN. With a SAN, storage or servers can be added or removed across multiple networks without impacting the availability and the accessibility. The SAN is optimized for “block I/O” which means a direct access to (virtual) disk sectors; the SAN performs SCSI commands, not file system-related commands. A key feature of SANs is their heterogeneous nature, with the ability to concurrently provide storage for servers running different operating systems. This feature is key for large environments needing to sustain high performance and availability.

## Internet SCSI

SANs rely on a very fast data transfer infrastructure which supports many devices linked over long distances. This need was initially satisfied by the Fibre Channel interface. The iSCSI (internet SCSI, or IP SCSI) specification defines a method to encapsulate the SCSI protocol commands within an IP packet. This allows SANs to be extended and to make use of conventional LAN or WAN topologies. iSCSI allows storage to be accessed over a LAN, allowing shared access to storage. iSCSI benefits from using existing, standard IP-based tools and services.

## A new layer of virtualization

The most recent storage virtualization features allow for the combining of the capacity of multiple storage controllers into a single resource with a single view of the storage resources. This abstraction layer between the physical storage devices and the users of those devices (host applications) provides the ability to hide the physical infrastructure from the application and the end user.

- At a block level with products such as IBM TotalStorage Volume Controller:  
Block Virtualization can be implemented using the SAN Volume Controller to allow the application systems to believe they have physically contiguous storage volumes allocated to them but in reality, the allocation is scattered across multiple devices. This allows the IT organization the flexibility to move application data from one physical device to another without affecting the application systems. The “virtual” disks that the application systems see appear to be unchanged. However, behind the scenes, the data may actually have been relocated to other physical devices.

- At a file level with products such as IBM TotalStorage SAN File System:  
File Virtualization can be implemented using the SAN File System to allow application systems to believe they have a locally attached file system, but in reality, the file system is shared across the Storage Area Network. This allows the administrator to designate policies to direct which type of storage each file is placed on, allowing some files to get higher Quality of Service (QoS) while other files get lesser service.

## Summary

The result of the focus on Open interfaces, combined with the modular components in the IBM storage solution, allows to offer flexible solutions in an heterogeneous storage infrastructure, with the tools to implement and manage this infrastructure with a global end to end view.

### 2.4.4 Network virtualization

There are business-critical application requirements to manage and utilize network resources more efficiently in regard to performance, resource usage, people cost, availability and security.

#### Network standardization

An important change in the past permitted to make this evolution: the network standardization:

- Ethernet has been chosen against SNA or Token Ring, and the evolution of the Internet has its responsibility part. The hardware components such as hubs and switches overtook MAU, bridges, and so on. This uniformization increased the research in the same way and we assist to the exponential progression of this technology, starting from 10mb/s up to 1000Mb/s now.
- TCP/IP has also been chosen as the default standard communication protocol; Novel IPX or Microsoft NetBEUI are now used in specific domains. Again the Internet helped a lot in this way. All associated subcomponents of this protocol such as DNS and DHCP have an important role in the identification of machines in the network.
- More recently, the capability to use fiber channels has increased the potential throughput and allows some new use of the network (with remote attached disks, and SAN; see “Storage virtualization” on page 33).

This network standardization increased the capability to communicate better inside the enterprise itself (between workstations, UNIX servers and mainframes for instance), but also make it easier between enterprises: Intranet and Extranet appeared.

## **The price of success**

The growth of the network utilization created new troubles in enterprises: Huge interconnections needed many wires, in machine rooms, in dedicated network rooms, and were difficult to manage. Often static, any changes needed a human intervention to plug or unplug a wire from a switch to another. To fix that, network boxes became more and more flexible and “smart” offering remote administration and permitting the creation of several networks on the same physical infrastructure: Virtual Networks and VLAN appeared.

Due to the partitioning evolution of servers (see “Server virtualization” on page 27), new network features were created inside boxes to create very high speed networks, as fast as a memory bus such as the HiperSockets available on zSeries. In xSeries Blade machines, the Blade Center itself hosts a switch module to create a LAN between Blades themselves. The network exit point from the Blade Center is unique. These evolutions allow internal communications by reducing the number of adapters and wires and increasing the management of the machine room.

More and more, new features were added to increase the network flexibility: today, network virtualization includes the ability to manage and control portions of a network that may even be shared among different enterprises, as individual or virtual networks, while maintaining isolation of traffic and resource utilization.

The automation of administration and maintenance tasks such as changing the associated VLAN of a machine is the key of the provisioning. Instrumentation of network resources and operations (SNMP, CIM, CLI, and so on) which can be abstracted across the server and networking devices are key enablers for on demand behavior.

## **Network and On Demand Operating Environment**

From an on demand perspective, network resources must be integrated into all the Service Level Management functions. Mapping the required network resources (including notions of bandwidth and priority) to the application/service being deployed is one of the key tasks that need to be automated. By virtualizing those network resources (server and core network), the ability to share them is enabled and management is simplified. The direction is for users to express their business goals via policy, and then to automate the process of translating those goals into resource-specific actions which can be coordinated to deliver the desired qualities of service.

It is important to note that much of the virtualization is at the platform layer, requiring support in the hypervisors — and in some cases, in the firmware — to enable sharing between different operating systems (in addition to the sharing that is provided in the operating system for functions such as Virtual LocalArea Network (VLAN), Quality of Service (QoS) and Virtual Private Network (VPN).



Also note that integrating the management of these networking resources in the context of the applications they support and the servers on which the applications run begins to reduce the complexity of managing servers and networks separately (for example, IBM Director support for Cisco and Nortel blades, and IBM Tivoli Provisioning Manager support for configuring server network resources such as adapters, IP addresses, VLANs, and so on).

The platform is also involved in provisioning/configuration, performance management, security and availability services. Functions such as intrusion detection filters may also be integrated in some operating systems. The use of policy to abstract the management of these services is also an element of the network virtualization support which is being enhanced to better integrate with the other resources (server/storage) and exploited by more of the service-level manager disciplines.

What we are driving toward is the notion of life cycle management of resources, where they are provisioned, monitored, and managed according to business goals using standards-based instrumentation and operations to reduce the complexity for customers in managing these disparate sets of resources from different vendors. Abstraction of physical interfaces and the support for sharing them across different virtual servers are the key concepts which network virtualization addresses.

## **2.4.5 Distributed systems**

In an On Demand Operating Environment, the convergence between the business view and the IT view is at the middleware level:

- ▶ On one side, the business application is hosted in the middleware.
- ▶ And on the other side, the middleware will use the virtualized infrastructure described previously.

The enhancements in the middleware demonstrate such convergence; for example: HTTP offers load balancing, WebSphere includes a concept of clusters, DB2® can offer Data Sharing among multiple servers. All these capabilities permit the deployment of an application in order to increase performance and scalability. The end users do not know which servers are used to handle their requests. This is a kind of virtualization.

Some years ago, two new technologies emerged: the Web Services and grid technologies. The former is used in classical Web applications to provide functions through services. The latter relies on the capability to provide machines' resources to handle a specific job. First used in deep computing and scientists domains, grid can now handle the needs of other sectors.

Both of these functions deliver a level of virtualization and abstraction of business processes that can allow businesses to offer more flexible and more dynamic access to their computer systems and services.

Apart from these technologies, specific middleware have integrated these virtualization concepts:

- ▶ IBM WebSphere in the XD version provides dynamic allocation features, including interesting features, such as Quality of Service (QoS) functions.
- ▶ IBM DB2 Information Integration will provide a uniform layer to access data, whatever the source is (databases, XML, queues, and so on...). It gives a virtual and common layer for data to increase the flexibility and make this data easier to access as an application design point of view.

This level of application virtualization offers significant business growth opportunities by allowing an organization to externalize key business processes in areas in which the organization can add significant value.

## Web Services

Web Services allows a computer program to dynamically locate a partner program that is an interface to a specific service. Web Services is a new model for using the Internet; it enables transactions initiated automatically by a program. The programs or services can be described, published, discovered, and invoked dynamically in a distributed computing environment. These services enable not just new business opportunities, but also new ways of using the Internet:

- ▶ Marketplaces
- ▶ Auctions
- ▶ Intelligent Agents

The services can be built on industry open standards, and exchange messages and data formatted with and described by the *eXtensible Mark-up Language* (XML).

Web Services has defined interfaces and technologies to allow an organization to describe the functionality (services) it wants to externalize. How and where they publish information about the service, and how other organizations discover services, connect to each other, and invoke services with appropriate security, reliability, and confidentiality, are all addressed by Web Services standards. XML defines a platform-independent way of representing data, making data integration easy and standardized. Web Services defines a platform-independent way of exchanging that data, so that process-level integration becomes much easier.

For more detailed and recent information about Web Services technologies and their use, refer to the IBM Redbook *Using Web Services for Business Integration*, SG24-6583.

## **Grid computing**

Grid computing is distributed computing based on open standards, enabling heterogeneous resources. Grids create a virtualized data center by linking servers, clients and storage to form virtual resource pools, which can be dynamically allocated. Grids enable virtual, collaborative organizations that share applications and data in an open, heterogeneous environment. Grid is a key enabling technology for e-business on demand.

Open Grid Services Architecture (OGSA) included together Web services standards, for example:

- ▶ WSDL (Web Services Description Language): an XML document for describing how to operate the service.
- ▶ SOAP (Simple Object Access Protocol): a means of messaging that connects the application and data between service provider and service requestor.

In OGSA, the term architecture denoted a well defined set of basic interfaces, which serves as the basis for constructing various systems; and the term open described extensibility, vendor neutrality and commitment to a community standardization process.

Grid is based on distributed computing. It means that the infrastructure is not centralized because their resources are geographically distributed. This implies some kind of communication between them. The network is the physical structure that interconnects the distributed resources. A network generally only provides the pipes to connect some component with each other. However, in order to create comprehensive and interesting communications, it is necessary to adopt some protocols and open standards.

Open standards and protocols provide the mechanisms for communication between components developed by different vendors. They also allow the developers to concentrate on the business logic rather than on programming communication routines.

The Globus Alliance and Globus Toolkit play an integral role in the advancement of grid technology:

- ▶ IBM is working very closely with the Globus Alliance, a multi-institutional research and development effort, to develop new architecture proposals and designs and to produce openly distributed reference implementations of architecture. IBM's contribution to The Globus Alliance includes sponsorship as well as actual design and software contributions.
- ▶ The Globus Toolkit is an integrated set of tools and software that facilitates the building of grids and the development of grid applications. It addresses issues of security, information discovery, resource management, data management, communication, fault detection and portability.

IBM Globus Toolkit will be, in the short term, an important coordination point and the glue for many virtualization technologies (see Virtualization Engine chapter).

Web Services interfaces have been used to enable integration of diverse distributed applications and middleware functions into end-to-end IT process in support of the business with *Web Services Resource Framework* (WS-RF). Web Services are extended to the heterogeneous distributed IT infrastructure, making possible the integration needed to view and manage it as a single compute resource.

WS-RF includes more and more grid functions, and we will assist in the convergence of a unique solution, as shown in Figure 2-5.

For more detailed and recent information about grid technologies and their use, refer to the IBM Redbook *Grid Computing with the IBM Grid Toolbox*, SG24-6332.

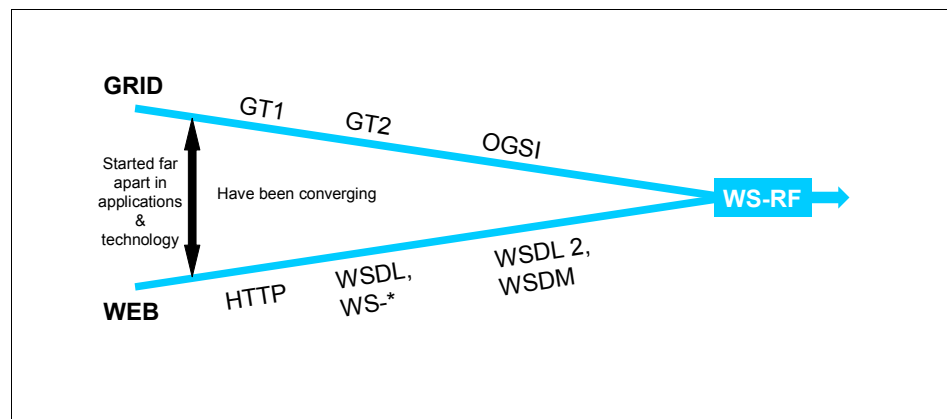


Figure 2-5 Web Services Resource Framework (WS-RF)

## 2.4.6 The IBM Virtualization Engine

Through the on demand initiative, it was important to present all the virtualization features with a single entry point. The IBM Virtualization Engine (VE) presents a new coherent and complete picture of the virtualization capabilities. Even if some of them are pretty new, others have been in existence for a long time already. This means that the virtualization concept is not a new theoretical fashion in marketing! The IBM Virtualization Engine provides the entry point for the virtualization capabilities, including hardware and software capabilities, as well as old and new capabilities.

The first release of the IBM Virtualization Engine was announced on August 17, 2004 and was made available on August 27, 2004. It is made up of multiple components and it is described in more detail in the next chapter.

- ▶ *Systems Technologies* consist of core technologies that will be bundled as standard elements in all IBM @server® lines over time. They will group components from servers, network, and storage virtualization, which we have described previously. They will include more components delivered through the hardware features.
- ▶ *Systems Services* offers two software suites to provide a new level of the infrastructure management.
  - *IBM Virtualization Engine Suite for Servers* is an integrated suite of virtualization software (including a common management console) that will be used to manage heterogeneous infrastructures.
  - *IBM Virtualization Engine Suite for Storage* consists of storage virtualization capabilities that will manage the two latest storage technologies which are SVC and SFS.





## The IBM Virtualization Engine

The IBM Virtualization Engine can help deliver to your company the promise of On Demand Business by simplifying systems management across operating systems, servers, and storage platforms; thereby helping to reduce the complexity, costs, and efforts associated with a heterogeneous IT environment. The IBM Virtualization Engine can reduce administrative and support costs by helping automate resource-tuning within and across systems; enabling more effective processor utilization; and improving availability by reducing the number of outages due to overutilized resources. Through its components, the IBM Virtualization Engine can also improve quality of service by identifying response time problems before they become critical and providing the optimum resources to the most important application when it needs it; and increasing utilization of servers by optimizing resources across servers.

The IBM Virtualization Engine provides a standardized virtualization foundation for the on demand architecture. It is common to, and can be managed and orchestrated across, all of IBM's server environments.

The IBM Virtualization Engine integrates a flexible mix of systems technologies and systems services, which better optimizes computing resources and integrates technology and business processes. These advantages help your organization respond in real time to market changes, opportunities, and customer demands.

In this chapter, we discuss the systems technologies and the systems service offerings that are part of the IBM Virtualization Engine, including the trends and directions anticipated.

## 3.1 Overview of the IBM Virtualization Engine

The IBM Virtualization Engine is a suite of systems services and technologies that provides a comprehensive approach to managing resources, servers, storage, operating systems, and networks as an integrated system, rather than as individual components. It reduces costs, improves storage utilization, and increases the amount of managed resources per administrator by simplifying the infrastructure and reducing management complexity.

The IBM Virtualization Engine is a key part of an On Demand Operating Environment and helps you optimally manage your current heterogeneous IT environment in a simpler, more efficient, and more responsive manner to meet the goals of your business. The IBM Virtualization Engine does not require “rip and replace” hardware and software upgrades. The IBM Virtualization Engine provides you with a high-level view of the computing resources that make up your enterprise; by doing so, it hides the details of the underlying diverse systems and hardware. It includes powerful tools for automating systems provisioning, workload management, storage virtualization, and partitioning.

Here are some of the key capabilities being provided by the IBM Virtualization Engine:

- ▶ It discovers and pools resources servers, storage, networks, both IBM and non-IBM, into a virtual environment.
- ▶ It shares and optimizes resources within a box and across the network.

The IBM Virtualization Engine consists of a set of system technologies and systems services which provides a consistent logical view of your cross-platform IT environment:

- ▶ Systems technologies are typically delivered with the hardware architecture. So, there are platform-specific capabilities delivered with the underlying hardware architecture and, typically, with the processor itself. A good example of that is logical partitioning capability. We have logical partitioning capability across all the series, as well as partitioning capability in storage. However, each one of the series delivers and implements logical partitioning in a unique way. Virtualization Engine will present each one of the unique delivery vehicles in a consistent manner.

Systems technologies are discussed in “Systems technologies” on page 47.

- ▶ System services span processor architectures, span operating system implementations, and are typically delivered to manage a distributed, heterogeneous environment. System services could support a single system, but they are designed and developed primarily to be used in a multi-system heterogeneous environment, such as the Enterprise Workload Manager (EWLM).



System services are discussed in more detail in “Systems services” on page 52.

The IBM Virtualization Engine is a combination of systems technologies and systems services, which provide the integrated virtualization technologies needed to support an On Demand Operating Environment.

IBM's virtualization technologies are built with open, industry standards, so customers can choose to implement all or any of them as their business requires. If other vendors participate in the open standards, they will be able to inter-operate with IBM's Virtualization Engine and on demand solutions.

Next, we discuss the VE components in more detail.

## 3.2 Systems technologies

The IBM Virtualization Engine Systems technologies are functions built into each IBM @server and TotalStorage platform that provide unique value for optimizing resources within a single system.

The IBM Virtualization Engine Systems Technologies include such functions as:

- ▶ Hypervisor™, which helps partition and move resources dynamically in multiplatform environments
- ▶ Virtual Ethernet, which allows the efficient prioritization of traffic on shared networks
- ▶ Virtual I/O, which allows flexible allocation, management, and sharing of physical resources (storage, adapters, and devices) among multiple partitions

The systems technologies also provide functions such as simultaneous multithreading, dynamic logical partitioning, and uncapped partitions for efficiently sharing server processor resources. Some of them are described in the following sections, organized by @server family.

### 3.2.1 Systems technologies for the zSeries family

Systems technologies for the zSeries family are designed to provide state-of-the-art capabilities and world-class support for workloads running within the zSeries architecture environment

Next we list some of these technologies.

- ▶ **Intelligent Resource Director (IRD):**  
IRD, introduced in 2000, extends the Parallel Sysplex architecture by helping you make sure that all the resources are being utilized by the right workloads, even if the workloads exist in different Logical Partitions. IRD uses facilities in z/OS Workload Manager (WLM), Parallel Sysplex, and PR/SM. IRD gives you the ability to move dynamically the resource (processor, channel) to where the workload is.
- ▶ **zSeries Application Assist Processor (zAAP):**  
zAAP is a Processing Unit (PU), similar to the System Assist Processor (SAP). They are designed to operate asynchronously with the general CPs to execute Java programming under control of the IBM Java Virtual Machine (JVM). IBM does not impose software charges on zAAP capacity. The zAAP optional assist feature allows customers to purchase additional processing power exclusively for Java application execution without affecting the total software pricing model or machine model designation.
- ▶ **HiperSockets:**  
HiperSockets provides internal “virtual” Local Area Networks which act like TCP/IP networks within the zSeries server. HiperSockets eliminates the need to utilize I/O subsystem operations and the need to traverse an external network connection to communicate between LPARs in the same zSeries server. HiperSockets provides the fastest communication between consolidated LPARs. This is an integrated zSeries Licensed Internal Code (LIC) function, which is coupled with supporting operating system device drivers (Linux, z/VM, VSE/ESA and z/OS virtual servers).
- ▶ **Workload Manager (WLM):**  
The workload manager is a component of z/OS. The purpose of WLM is to balance the available system resources to meet the demands of S/390 subsystems work managers such as CICS, Batch, TSO, UNIX System Services, and Webserver, in response to incoming work requests. The workload management services enable z/OS to cooperate with subsystem work managers to achieve installation-defined goals for work, to distribute work across a sysplex, to manage servers and to provide meaningful feedback on how well workload management has achieved those goals. Most of the WLM for z/OS concepts have been used for the new EWLM, described in “IBM Enterprise Workload Manager” on page 52.
- ▶ **Dynamic Logical Partitioning (LPAR),** described in Section , “Logical partitioning” on page 29, Parallel sysplex clustering, described in “With the Parallel Sysplex” on page 31, Virtual LANS (VLANs) are other virtualization technologies available into the IBM zSeries.

### 3.2.2 Systems technologies integrated into the pSeries

Virtualization systems technologies integrated into the pSeries include these:

- ▶ **POWER™ Hypervisor:**

The POWER Hypervisor is a piece of firmware on the pSeries servers which is the underlying control mechanism that resides below the operating systems, but above the hardware layer. This firmware performs the initialization and configuration of the POWER5 processor, as well as the virtualization support required to run up to 254 partitions concurrently. With support for dynamic resource movement across multiple environments, customers can move processors, memory, and I/O between partitions on the system, as workloads are moved between the partitions.

The POWER Hypervisor supports many advanced functions, such as sharing of processors, virtual I/O, high-speed communications between partitions using a Virtual LAN, concurrent maintenance, and also allows for multiple operating systems to run on the single system. Currently, the AIX 5L™, Linux and i5/OS™ operating systems are supported.

- ▶ **Network Installation Management (NIM):**

NIM provides remote installation of the operating system, manages software updates, and can be configured to install and update third party applications. NIM is provided as part of AIX.

- ▶ **Micro-Partitioning:**

Micro-Partitioning provides the ability to share processors among partitions in the system. Micro-Partitioning is the mapping of virtual processors to physical processors. Virtual processors (not physical) are assigned to the partitions. With the POWER Hypervisor, an entitlement or percentage of processor usage is granted to the shared partitions. The granularity of processor usage is ten percent. By dividing up processor usage in this manner, a system can have multiple partitions sharing the same physical processor.

Micro-Partitioning allows several operating systems to share the physical processor resources in a time-sliced manner. AIX 5L version 5.3 is the first version of AIX to support Micro-Partitioning.

- ▶ **Partition Load Manager (PLM):**

The Partition Load Manager is a resource manager that assigns and moves resources based on defined policies and utilization of the resources. PLM manages memory, both dedicated processors and partitions using Micro-Partitioning technology to readjust the resources. PLM, however, has no knowledge about the importance of a workload running in the partitions and cannot readjust priority based on the changes of types of workloads.

- ▶ Virtual Ethernet, Dynamic LPAR, Shared Ethernet Adapter, Virtual SCSI, Virtual I/O (I/O appliance, VLAN, Virtual SCSI Server) are other virtualization technologies available into the IBM pSeries.

### 3.2.3 Systems technologies integrated into the iSeries

The IBM @server i5 has been the industry's first POWER5 processor-based server. The IBM Virtualization Engine technologies enabled by the Power Hypervisor include:

- ▶ Dynamic Logical Partitioning (LPAR):  
This capability allows the IBM @server i5 system resources to be grouped into logically separate systems within the same physical footprint. With capabilities such as *uncapped partitions*, the processors on the i5 server can be shared between partitions based on business needs (defined via policies).
- ▶ Virtual Ethernet:  
This capability allows the partitions to communicate within a high-speed virtualized network, yet all have outside access via routing through one set of physical I/O devices
- ▶ Virtual I/O:  
I/O resources such as disk, tape, and CD-ROM can be shared between multiple operating system partitions and IXS/IXA hardware.
- ▶ Capacity on Demand:  
System resources such as processors and memory are made available on an as-needed basis<sup>1</sup>. Once activated, the resources can be used temporarily (On/Off CoD) or permanently (CUoD) when the business need arises.
- ▶ Simultaneous Multithreading:  
Applications can increase the overall resource utilization by virtualizing multiple physical CPUs through the use of multithreading. Simultaneous Multi-threading allows two instruction paths to share access to the POWER5 execution units on every clock cycle. Each instruction path is abstracted, by the operating system, so that each POWER5 virtual processor will appear as two logical POWER5 processors, one for each instruction path. With no modifications to the applications, simultaneous multi-threading will enable SMP-scalable applications to benefit from significantly increased system level performance (IBM relative performance, rPerf, projections show up to 30% improvement in utilization and throughput).
- ▶ Multiple Operating Systems:  
Multiple Operating Systems, deployed simultaneously on separate logical partitions, are supported on an IBM i5 server. The operating systems

supported by the i5 include AIX 5, i5/OS, and Linux for POWER. The i5 also supports the Microsoft Windows 2003 Server operating system, via the Integrated xSeries Server/Integrated xSeries Adapter (requires at least one partition of i5/OS).

- ▶ Micropartitioning enables you to allocate less than a full processor to a logical partition.

### 3.2.4 Systems technologies integrated into xSeries and BladeCenter

Virtualization systems technologies integrated into the xSeries and BladeCenter include:

- ▶ VMWare:

The VMware virtualization layer (a non-IBM product) brings hardware virtualization, pioneered on IBM VM/370 and other mainframe environments, to the standard Intel server platform. The virtualization layer provides an idealized physical machine that is isolated from other Virtual Machines on the system. It provides the virtual devices that map to shares of specific physical devices. These devices include virtualized CPU, memory, I/O buses, network interfaces, storage adapters and devices, human interface devices, BIOS and others. Each VM runs its own operating system and applications. They cannot talk to each other or leak data, other than through networking mechanisms similar to those used to connect separate physical machines.

- ▶ Virtual Machine Manager:

IBM has developed a plug-in to IBM Director called Virtual Machine Manager (VMM) that extends IBM Director to include virtualization awareness. Specifically, VMM supports both VMware deployments and Microsoft Virtual Server deployments, and it provides a detailed view of physical platforms correlated to the virtual environment that each of those supports. VMM can detect a hardware alert via the base IBM Director software and migrate one or more virtual machines off that server via VMotion onto another server that is operating normally with sufficient resources to support additional workload.

- ▶ Capacity manager:

Capacity Manager is a system management tool which is part of IBM Director and which can help you to measure the potential bottlenecks of multiple IBM Director systems. You can use this tool to forecast performance degradation of a server and its subsystems. You can plan for an appropriate action to overcome the bottleneck well in advance, so as to prevent overall performance degradation. Capacity Manager for IBM Director 4.20 is supported on all IBM Director Agents running under Linux, NetWare or Windows operating systems.

- Integrated shared infrastructure for Blades:

For example, in a BladeCenter, all service is from the front and rear of BladeCenter. There is no need to slide the chassis out of the rack and remove the top cover for service. Also, numerous cables are eliminated, reducing both cabling cost and service/administrator time.

## 3.3 Systems services

The IBM Virtualization Engine Systems Services are the “fabric” that bind the various platform technologies together so they can operate cohesively as a single entity. The IBM Virtualization Engine systems services consists of two suites: IBM Virtualization Engine Suite for Servers and the IBM Virtualization Engine Suite for Storage. The IBM Virtualization Engine System Services for @server uses IBM's Integrated Solutions Console as its user-friendly, dashboard-like view of the status of your environment. The IBM Integrated Solutions Console is a portal-based console that lets you keep track of the systems and resources in your environment by allowing common systems administration across product boundaries.

Increasingly, as organizations become more virtualized and systems are acquired more dynamically, the ability to configure to get the most out of existing and new systems requires a common set of management tools.

The IBM Virtualization Engine goes a long way toward addressing this requirement for servers and storage by abstracting out the differences in implementation and providing a consistent and effective way to deal with different hardware and operating system implementations.

### 3.3.1 Suite for Servers

The IBM Virtualization Engine Suite for Servers is a pretested, multi-platform offering of systems services which can help provide virtualization and management of resources across a variety of operating system environments, both in an individual server and across an enterprise. It consists of the following key components.

#### **IBM Enterprise Workload Manager**

IBM Enterprise Workload Manager will enable you to automatically monitor and manage heterogeneous workloads across an IT infrastructure to better achieve predefined business goals for end-user performance. It provides end-to-end resource optimization and load balancing of IT resources in heterogeneous, multi-tier application and server environments. See Figure 3-1.

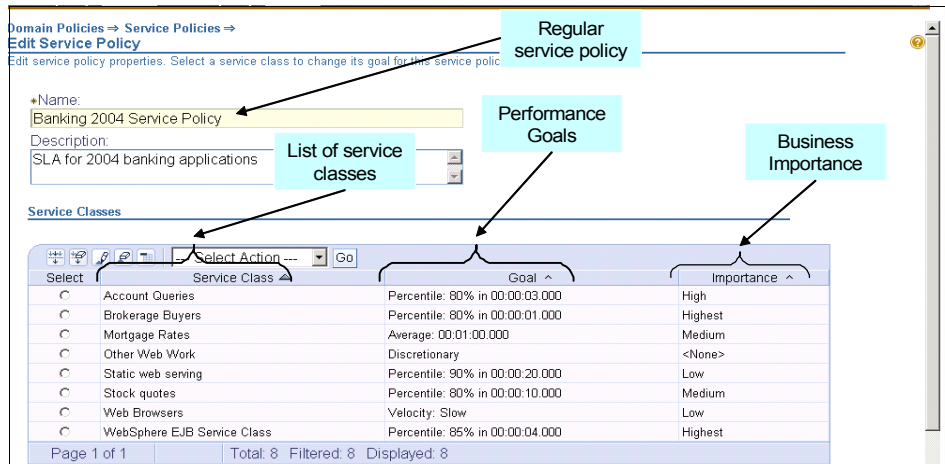


Figure 3-1 Enterprise workload monitoring

The goal of EWLM is to provide the ability to identify work requests based on service class definitions, track performance of those requests across server and subsystem boundaries, and manage the underlying physical and network resources to achieve specified performance goals for each service class. Figure 3-1 describes some of the EWLM components as seen through the EWLM Control Center.

EWLM management capability will be delivered in stages across several releases. In this first release, it assists network routers in distributing work to those servers best able to achieve service class goals for the specific category of work being managed as described in Chapter 17, “How to balance workloads in the network” on page 255.

By bringing this self-tuning technology to the set of servers and routers, Enterprise Workload Manager can help enable greater levels of performance management for distributed systems. Its extremely flexible management reach can encompass your infrastructure or be deployed to handle a critical subset of your applications.

EWLM offers monitoring capabilities to allow analysis of activities across your enterprise, and manage resource allocation and utilization based on business requirements.

EWLM technology provides the functionality to manage the heterogeneous servers that are defined within what is called an Enterprise Workload Management domain. Each Enterprise Workload Management domain can have hundreds or even thousands of server resources within it.

EWLM essentially consists of four pieces: a domain policy, a domain manager, a managed server, and an administrative interface. They are detailed in the following sections.

### ***Domain policy***

These are the policy-based service level objectives for the workload that will run on the servers within the Enterprise Workload Management domain. The objectives are defined by establishing service policies and associating the workloads to a policy. These definitions are entered using the administrative user interface and are sent to the domain manager where they are stored in XML format.

While only one policy can be active at any time, you can predefine different policy settings, changing them dynamically as your business needs dictate, such as having different policies for prime shift, off-shift, or weekend processing.

### ***Domain manager***

There is one physical domain manager for each Enterprise Workload Management domain of servers that are being managed. The domain manager is responsible for storing the domain and service policies. In addition, it dynamically manages the server topology within the Enterprise Workload Management domain and maintains it, along with the state of each server. It also handles communication with the administrative user interface as well as with all servers in the Enterprise Workload Management domain. The domain manager is currently supported on AIX, Linux, OS/400®, and Windows. The EWLM Domain Manager for z/OS is expected to be enabled in a future release allowing customers to monitor and manage EWLM from their z/OS system.

Finally, it aggregates the performance data collected by each of the management servers into a comprehensive, end-to-end view of overall service class and server performance.

### ***Managed server***

This is code that runs on each server within the Enterprise Workload Management domain being managed. Each server contains an implementation of the code, which is based on the Application Response Measurement (ARM) standard from the Open Group<sup>1</sup>, either delivered with the operating system or installed along with Enterprise Workload Management, depending upon the platform. The managed server function is currently supported on AIX, OS/400, Solaris, Windows, and z/OS<sup>2</sup>.

---

<sup>1</sup> For more detailed information, please look at the following Web site:

<http://www.opengroup.org/tech/management/arm/>

<sup>2</sup> Starting with z/OS V1R6 (or z/OSe V1R6) z/OS includes EWLM ARM support.



This Enterprise Workload Management ARM implementation interfaces with the operating system as well as with instrumented software running on the server, such as:

- ▶ DB2 Universal Database™, including z/OS support for DB2
- ▶ WebSphere Application Server
- ▶ Web serving plug-ins.
- ▶ IBM Dynamic Infrastructure Enterprise Edition for mySAP Business Suite

**Note:** The WebSphere Application Server ships the following Web application servers, which are instrumented by default because they use the ARM instrumentation framework provided by WebSphere Application Server.

These servers include, but are not limited to:

- ▶ Apache HTTP Server
- ▶ Covalent Enterprise Ready Server
- ▶ Covalent Enterprise Ready Server FastStart
- ▶ Covalent FastStart Server
- ▶ IBM HTTP Server
- ▶ Lotus® Domino® Enterprise Server
- ▶ Microsoft Internet Information Services (IIS)
- ▶ Sun ONE Web Server, Enterprise Edition

In addition, IBM has ARM-instrumented the following non-WebSphere Application Server Web serving plug-ins:

- ▶ Independent plug-in for Apache HTTP Server 2.0.47 or later
- ▶ Independent plug-in for IBM HTTP Server 2.0.47 or later
- ▶ IBM HTTP Server for iSeries
- ▶ EWLM-supplied plug-in for Internet Information Services (IIS) 5.0 and 6.0. (Use this plug-in for static Web page serving only. It is not recommended for use in production environments. Deployment of this IIS plug-in for other than static Web page serving can result in inaccurate ARM transaction event reporting, and subsequently inaccurate EWLM interpretation of transaction)

Information is provided by the ARM instrumented interfaces, which allow the tracking of workload that is flowing through the servers in the Enterprise Workload Management domain.

Information that is gathered is sent to the domain manager to be summarized and analyzed. This information gives insight into how the server and middleware is running from a utilization standpoint, as well as how well the pieces of the workload are being serviced from a response time and utilization standpoint. The information is then available to be accessed from the administrative user interface.

### ***Administrative user interface***

This is the point from which the Enterprise Workload Management domain is controlled. The domain policy and service policies can be entered from this Web based interface. The interface is also the tool used to control when domain policies and service policies are to be implemented. This interface can be used to instruct the domain manager to implement certain domain policies and service policies for the environment. The domain manager communicates with the managed server components running on each of the servers to obtain performance information.

By being able to know the utilization of a server, plus the middleware component associated with the server, plus the response time of each of these components and associating all of this knowledge with the service level objective for the workload running on these servers, it is possible to determine if the service level objectives are being met for the workload as well as whether the server can be utilized at a higher rate while not jeopardizing the SLA for the workload.

More detailed information about EWLM can be found in the redbook, *IBM Enterprise Workload Management*, SG24-6350.

### **IBM Director Multiplatform**

IBM Director Multiplatform is the component of the IBM Virtualization Engine that provides a simplified view of IBM @server hardware from a centralized server, eliminating platform unique management approaches for hardware and virtual partition management. It provides a comprehensive standards based systems management solution for heterogeneous environments, such as support of IBM @server BladeCenter, iSeries, pSeries, and xSeries servers, network and storage. IBM plans to add similar functionality in future releases of IBM Director Multiplatform for supported platforms, including IBM @server zSeries servers.

IBM Director Multiplatform's core infrastructure is designed to provide a single point of control for managing up to 5,000 systems.

The Director Multiplatform product structure consists of three components:

- ▶ The Management Server is the main component or aggregation point for managing the agents.
- ▶ The Agent provides management data and function to the Management Server. Depending on the operating system and hardware platform, the Agent capabilities will vary.
- ▶ The Console provides a graphical user interface (GUI) with a consistent look and feel for all the servers and devices maintained by the Management Server.

By leveraging industry standards, Director Multiplatform provides an extensible architecture to enable solutions for easy integration with other systems management tools and applications, such as Tivoli Enterprise™, Tivoli NetView®, HP OpenView, Microsoft SMS, CA Unicenter, BMC Patrol, and NetIQ.

Using the collection of IBM Director Multiplatform tools, many of the administrator's manual tasks can be automated to proactively and remotely manage systems such as discovery, event logs and action plans, file transfer, inventory collection, process management, resource monitors and thresholds, and more. The predictive and proactive capabilities associated with alerting and real-time system diagnostics help maximize server uptime and reduce service downtime costs. Finally, and perhaps most importantly, the product's ability to support cross-platform IBM and non-IBM systems (Intel-based) means customers can protect existing infrastructure investments and can manage heterogeneous environments.

For xSeries servers and BladeCenter, Director Multiplatform provides additional exploitation of the hardware through inventory and alerts, asset tracking, diagnostic monitoring, capacity management and troubleshooting.

## **Virtualization Engine Console**

The Virtualization Engine Console is a Web-based console that provides two functions:

- ▶ A consolidated view of enabled enterprise resources via a health function
- ▶ A launch point for the Virtualization Engine system services, as well as other Web based consoles

The Virtualization Engine Console health function allows customers to see system resources across platforms. This is done using input from any of the following IBM products: IBM Director Multiplatform, Cluster Systems Management, iSeries Navigator, and IBM Tivoli Monitoring.

The Virtualization Engine Console is based on the IBM Integrated Solutions Console. This provides a common framework on which other consoles can plug in using a portal approach. This technology is one that IBM is extending across many of its product lines.

Figure 3-2 shows an example of the main page of the Virtualization Engine Console. The left navigator allows you to move between the health function and the launch pad function. You can also see a “dashboard”, which you can configure; it allows you to view the areas of your environment most important to you. By moving the mouse over any of the gauges in the dashboard, you can obtain more detailed information.

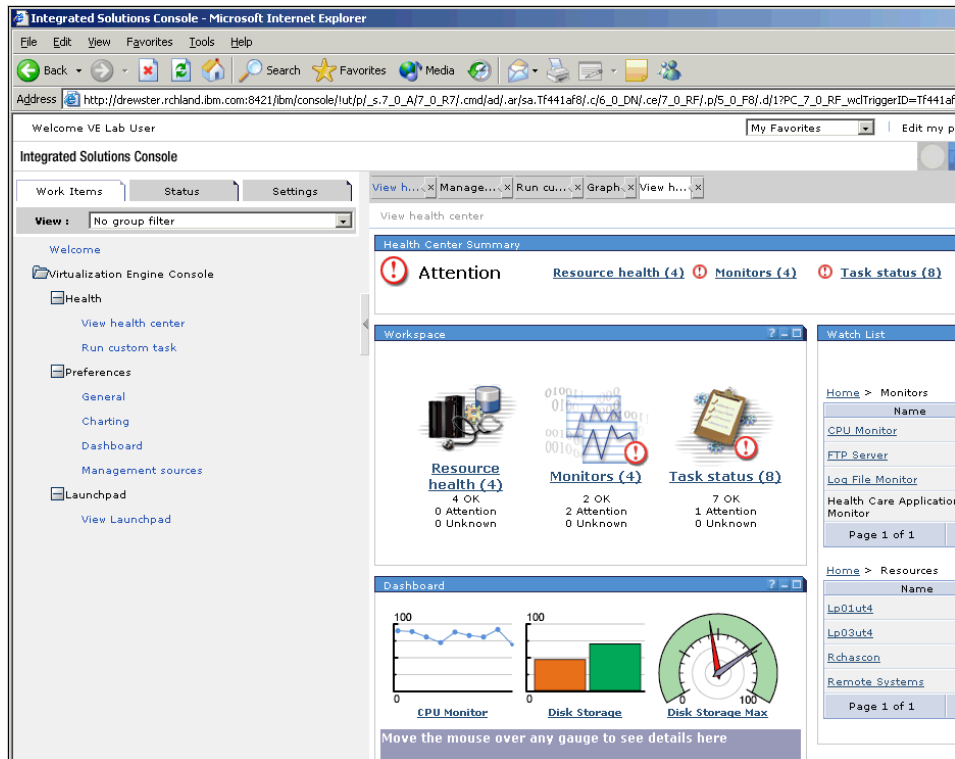


Figure 3-2 The Virtualization Engine console

Platforms currently supported by the IBM Virtualization Engine which can use the Virtualization Engine console to monitor the health of their environment are:

- ▶ IBM @server iSeries, pSeries, xSeries
- ▶ AIX 5L, i5/OS
- ▶ Microsoft Windows or Sun Solaris
- ▶ IBM TotalStorage

For specific component and platform availability information, refer to the IBM Virtualization Engine announcement.

## IBM Tivoli Provisioning Manager

The primary value of the systems provisioning component is to increase the average utilization of resources, thus reducing the number of required resources and the associated management costs. The IBM Tivoli Provisioning Manager does this by automating manual tasks of provisioning and configuring servers, operating systems, middleware, applications, storage and network devices. The IBM Tivoli Provisioning Manager works together with a set of thoroughly

tested and documented “best practices” workflows that support typical IBM @server and IBM TotalStorage topologies. These workflows can automate your unique data center processes, including installing, configuring, and deploying servers, operating systems, middleware, applications, storage, and network devices.

In addition to these @server workflows, users can also download and use any workflows from Tivoli's Orchestration and Provisioning Automatic Library (OPAL) Web page, or create their own customized workflows to implement their company's data center best practices and procedures. These procedures can then be automated and executed in a consistent, error-free manner. In fact, using these automation workflows, IBM Tivoli Provisioning Manager has the ability to provision and deploy a server (from bare metal to full production) with the single push of a button, reducing the time needed to deploy a new server from days to minutes.

One of the key functions of IBM Tivoli Provisioning Manager is platform provisioning. This is the ability to create and manage OS Containers. OS Container is a collection of resources (processors, memory, I/O, storage) configured to host an OS Standard Server, Configured Blade Server, LPAR, VMware partition, VM ready to run Linux on zSeries, and so on. It also includes the creation and management of operating system images to load in containers. Finally, platform provisioning includes the ability to install the operating system in an OS Container.

IBM Tivoli Provisioning Manager supports today the following environments:

- ▶ IBM @server xSeries and pSeries servers, including pristine installations
- ▶ IBM @server zSeries, that supports the cloning of manually installed Linux system images into automatically provisioned virtual machines under z/VM
- ▶ IBM @server iSeries, that supports provisioning of both Linux on POWER partitions
- ▶ Windows on Integrated xSeries Adapters and Servers.

Working with the IBM Tivoli Intelligent ThinkDynamic Orchestration (ITITO), you can extend the provisioning functionality to automate and orchestrate the provisioning of your IT resources, providing the appropriate resources on demand to your business applications.

IBM Tivoli Provisioning Manager can provide:

- ▶ Implementation of automation workflows
- ▶ Consolidation of testing center environments
- ▶ Application server support
- ▶ Storage capacity provisioning

## IBM Grid Toolbox Multiplatform

The IBM Grid Toolbox Multiplatform is a comprehensive, integrated toolkit for creating and hosting grid services. This product includes material developed by the Globus Alliance<sup>3</sup>, as well as a set of APIs and development tools to create and deploy new grid services and grid applications.

IBM Grid Toolbox Multiplatform V3 is designed to exploit the existing virtualization on IBM @server and TotalStorage families. It provides the ability for logical partitions, blades, and virtual machines to manage work across a heterogeneous, distributed environment. These programs can be responsible for managing specific tasks on the servers, for creating/deploying new services, and for participating in distributed system communication and management.

The IBM Grid Toolbox Multiplatform includes a limited integrated hosting and development environment capable of running grid services and sharing them with other grid participants, such as grid service providers and grid service consumers. It also provides a set of tools to manage and administer grid services and the grid hosting environment, including a Web-based interface, the grid Services Manager.

Currently, the toolbox supports the IBM AIX operating system on IBM @server pSeries and iSeries servers, i5/OS on IBM @server iSeries, and the Linux operating system on IBM @server xSeries, @server pSeries, IBM @server iSeries and IBM @server zSeries servers.

Derived from the common runtime used for the IBM Grid Toolbox Multiplatform, the IBM Virtualization Engine supports industry open standards that support the inter operability of heterogeneous IT resources. Standards and technologies such as the Open Grid Services Architecture (OGSA), Linux, XML, WSDL, and SOAP give IT solution providers a more cost-effective way to provide solutions for heterogeneous environments, while increasing the compatibility between components within an infrastructure. Support for these standards and technologies are planned across all IBM @server platforms.

## IBM Dynamic Infrastructure Enterprise Edition for mySAP

The IBM Dynamic Infrastructure Enterprise Edition for mySAP Business Suite is a feature of the IBM Virtualization Engine Suite for servers (available since January 2005). It can provide a resilient information technology (IT) environment that enables growth and capacity management in the SAP environment. This solution implements a SAP adaptive computing concept with additional functions around dynamic provisioning of resources for SAP applications.

---

<sup>3</sup> <http://www.globus.org/>

Dynamic Infrastructure is an IBM on demand solution for a heterogeneous environment that can enable you to run SAP environments more efficiently by dynamic provisioning and de provisioning of SAP systems across IBM platforms. It offers automated, policy-based, end-to-end management of resources, performance, availability, and security. It also offers accounting and metering information about resource consumption across heterogeneous systems, driven by service-level agreements (SLA). IBM Dynamic Infrastructure can reduce the high management effort for growing SAP landscapes, speed up the deployment of new SAP systems, and improve the utilization of systems, thus helping to lower the total cost of ownership (TCO). The combination of IBM's virtualization and provisioning technologies with this very cost-effective platform gives you a modern and competitive solution to support your growing SAP landscapes.

IBM Dynamic Infrastructure requires one or two management servers that run the management infrastructure components. Depending on the platform on which the provisioned SAP systems (that is, the managed systems) will run, the management servers run on:

- IBM pSeries (AIX 5L), or
- IBM xSeries (Microsoft Windows, Linux)

For the provisioned SAP systems, IBM Dynamic Infrastructure Enterprise Edition for mySAP Business Suite supports the following platforms and functions:

- ▶ On IBM iSeries, pSeries (both AIX 5L), and IBM zSeries (Linux):
  - Automatic or dynamic (SLA based) provisioning or de provisioning of SAP application servers
  - Metering, internal accounting, and external billing (via a separate billing engine)
- ▶ On xSeries and BladeCenter (Linux):
  - Start, stop, and relocate of both application and database server
  - High availability as an integral part of the solution

### **3.3.2 Suite for Storage**

The IBM Virtualization Engine Suite for Storage can help provide the capability to virtualize and manage storage devices and files. The key separately delivered products are described in the following sections.

#### **IBM TotalStorage Productivity Center**

The IBM TotalStorage Productivity Center is the first offering to be delivered as part of the IBM TotalStorage Open Software Family. The IBM TotalStorage Productivity Center is a component of the VE Suite for Storage that provides an open storage infrastructure management solution designed to help reduce the

effort of managing complex storage infrastructures and help improve storage capacity utilization and administrative efficiency. It simplifies the management of traditional and virtualized SAN environments, and is designed to enable an agile storage infrastructure that can respond to on demand storage needs.

The IBM TotalStorage Productivity Center is comprised of a user interface designed for ease of use, in addition to these components:

- ▶ IBM Tivoli Storage Resource Manager
- ▶ IBM Tivoli SAN Manager
- ▶ IBM TotalStorage Multiple Device Manager

The IBM TotalStorage Productivity Center is designed to help improve:

- ▶ Administrator efficiency
- ▶ Capacity utilization
- ▶ SAN performance
- ▶ Application availability

With IBM TotalStorage Productivity Center, you can:

- ▶ Monitor and track the performance of SAN-attached Storage Management Initiative Specification (SMI-S)-compliant storage devices
- ▶ Manage the capacity utilization and availability of file systems and databases
- ▶ Monitor, manage, and control (zone) SAN fabric components
- ▶ Manage advanced storage replication services (Peer-to-Peer Remote Copy and FlashCopy®)
- ▶ Automate capacity provisioning to help improve application availability

## **IBM TotalStorage SAN Volume Controller**

IBM TotalStorage SAN Volume Controller enables storage virtualization and can help increase the utilization of existing capacity. It is designed to increase the flexibility of your storage infrastructure by enabling changes to the physical storage with minimal or no disruption to applications through the function of pools of virtual volumes that can span several physical storage resources.

Now with expanded support for many non-IBM storage systems, the IBM TotalStorage SAN Volume Controller can enable a tiered storage environment to better allow you to match the cost of the storage to the value of your data. It is designed to allow you to centrally manage multiple storage systems to help enhance productivity, and combine the capacity from multiple disk storage systems into a single storage pool to help increase utilization. It also allows you to apply advanced copy services across storage systems from many different vendors, to help further simplify operations, by optimizing how storage resources are used, viewed, and centrally managed.



## **IBM TotalStorage SAN File System**

The IBM TotalStorage SAN File System (based on IBM Storage Tank™ technology) provides a network-based heterogeneous file system for data sharing and policy-based storage management in an open environment. It is designed to help reduce the complexity of managing files within SANs. The IBM TotalStorage SAN File System provides high availability; it is designed to provide a network-based heterogeneous file system for data sharing and centralized policy-based storage management in an open environment. It allows administrators to dynamically manage the resources based on user-defined policies.

IBM TotalStorage SAN File System is designed to enable host systems to plug-in to a common SAN-wide file structure. With the IBM TotalStorage SAN File System, files and file systems are no longer managed by individual computers; instead, they are viewed and managed as a centralized IT resource with a single point of administrative control.

The IBM TotalStorage SAN File System provides a common file system for UNIX, Windows and Linux servers, with a single global namespace to help provide data sharing across servers. It is a highly scalable solution supporting both very large files and very large numbers of files without the limitations normally associated with Network File System (NFS) or Common Internet File System (CIFS) implementations.

The IBM TotalStorage SAN File System is designed to help lower the cost of storage management and enhance productivity by providing centralized and policy-based storage and data management for supported heterogeneous server, operating system and storage platforms.

### **3.3.3 System services summary**

The IBM Virtualization Engine systems services are tested together and integrated to support the delivery of an On Demand Operating Environment. This is implemented in such a way that nonparticipating environments can be managed by and participate in On Demand Business, but do not have to make major changes to benefit from the IBM Virtualization Engine systems services and technologies.

In the IBM Virtualization Engine Suite for Servers, IBM Director Multiplatform centrally manages and self-monitors the hardware resources in your diverse IT environment in a homogeneous way. IBM Director Multiplatform works in conjunction with the Virtualization Engine Console to provide the common user interface into your distributed systems and with the IBM Tivoli Provisioning Manager to automatically provision your environment with the necessary applications and storage. IBM Enterprise Workload Manager autonomously optimizes and tunes your network according to business goals to improve the performance of your infrastructure.

In the IBM Virtualization Engine Suite for Storage, IBM TotalStorage Productivity Center with its IBM Tivoli Storage Resource Manager, IBM Tivoli SAN Manager, and IBM TotalStorage Multiple Device Manager, autonomously monitors the storage and devices in your cross-platform distributed environment. It automatically detects problems and provides self-healing capabilities through triggered alerts. IBM TotalStorage SAN Volume Controller centralizes SAN management and IBM TotalStorage SAN File System supports self-monitoring and policy-based automation in placing files. An On demand environment can provide an IT infrastructure that will help your business better align its IT operations with its business strategy.

Figure 3-3 describes how different components from the IBM Virtualization Technologies and System Services can work together, dynamically, when an event (in this example: a blade is introduced) happens inside the infrastructure.

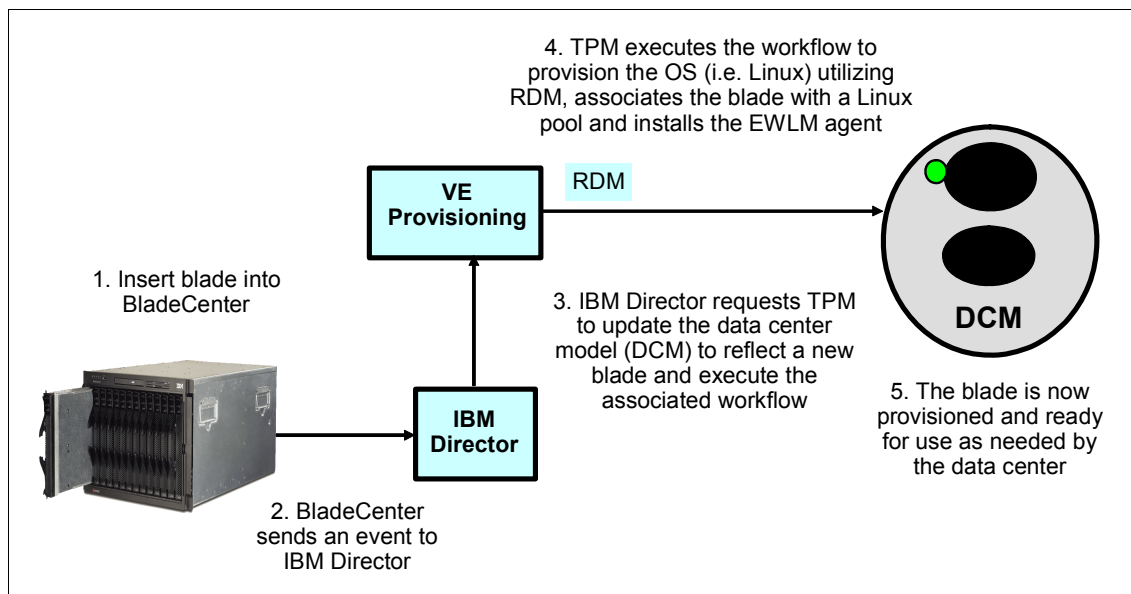


Figure 3-3 Virtualization Engine integration

In the On Demand Operating Environment, the Virtualization Engine's services and technologies are leveraged to provide a virtualized, flexible, and dynamic infrastructure that can support On Demand Business services. For organizations that are not ready for a full on demand environment, the Virtualization Engine services and technologies have their own, combined value proposition when running on IBM and heterogeneous systems.

## 3.4 Summary

In this chapter we examined the key role of the IBM Virtualization Engine (VE) and discussed its functions and services. We also looked at how the Virtualization Engine will be leveraged by the On Demand Operating Environment, as well as how they improve and deliver on the virtualization of key infrastructure components.

Having the ability to efficiently manage the IT infrastructure is a journey, not a major re-engineering project, requiring massive rip-and-replace steps. There is a stepped approach that most successful companies follow, as shown next in Figure 3-4.

The first thing that needs to be done is to simplify the environment by consolidating like systems onto fewer, more manageable resources. We have already supplied the systems technology, together with products such as the IBM Director Multiplatform to allow customers virtualize like resources. These are resources on homogenous systems, storage and networks; the virtualization system technologies as discussed previously in this chapter. This has been done time and time again by our customers. However, this is just the beginning.

Following that step, automation is key to manage those resources and being able to add or move resources as required by the business. Allowing business needs to drive resources will dictate how well the business performs. We have also began to deliver on the second stage of virtualization through the System Services to allow you to virtualize unlike system resources, application based grids and network virtualization.

As the evolution of the IBM Virtualization Engine continues, more platforms will be supported by the various Suite for Server and Suite for Storage products.

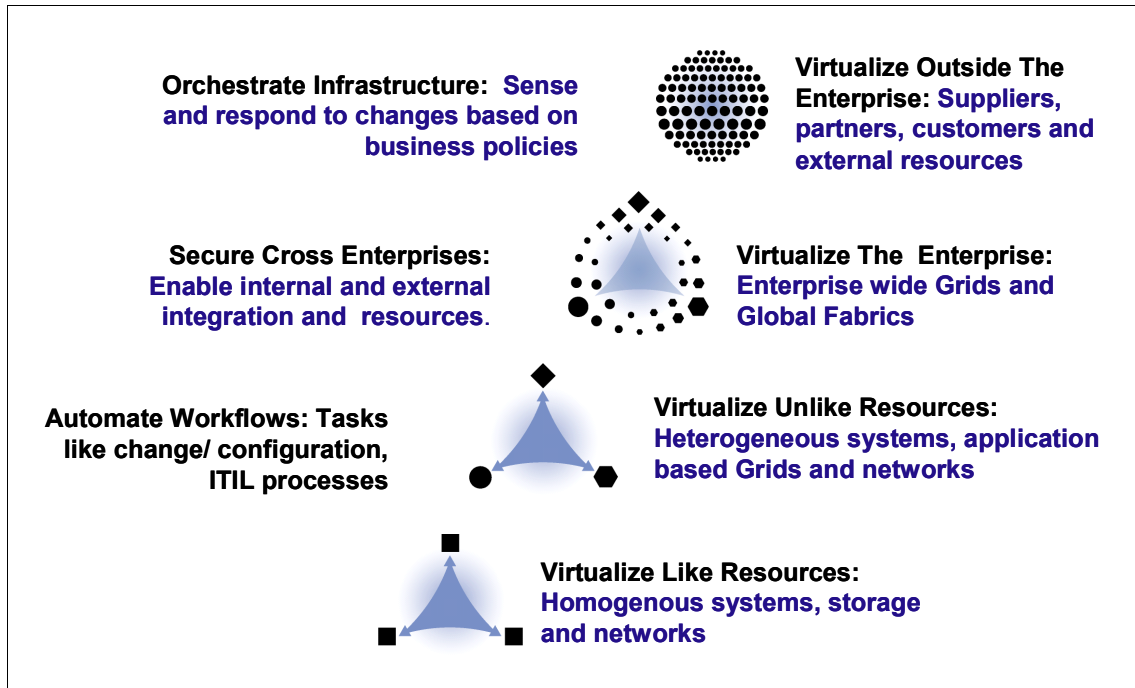


Figure 3-4 Stages of virtualization solutions



## Part 2

# How to's for managing the Infrastructure

Most enterprises evaluating the On Demand Operating Environment will consider a “start small” approach focused on an area or issue that addresses their current business imperatives or pain points. This is consistent with the vision of the On Demand Operating Environment, based on its modular technology framework and iterative methodology. It also allows enterprises to leverage their existing systems and skill base as they enhance their functions and capabilities.


In this part of the book, we describe several approaches for getting started with deploying an On Demand Operating Environment.

The different chapters of this part discuss some of the potential starting points for adopting infrastructure management components of an On Demand Operating Environment. Each approach is based on real-world business imperatives that enterprises face and uses elements of the technology framework described in the previous part of this book.

Each approach is discussed briefly with a strategy presentation, the description, and the role of the components that can be part of the solution.

In this redbook we do not intend to give an exhaustive list of the problems that may be solved using the On Demand Operating Environment components. This book is a first iteration, going through the most common infrastructure problems.

Some of these approaches are detailed through scenarios, and discussed in Part 3, “Infrastructure management: Detailed scenarios” on page 165.



## **How to secure access and control of information, resources, and applications**

Information in data stores such as databases, transactional systems, file systems, and even Web Services is becoming increasingly accessible within enterprises as well as from the external world. The same applies to enterprise operational entities such as applications, business process workflow, operating system processes, and traditional batch jobs.

The IT challenge is to secure and control access to information and operational entities automatically in real time. This requires synchronization of identity across the infrastructure as well as a consistent set of policies regarding access control. The goal is to provide controlled, yet pervasive access to information and applications as it is needed, based on the demands of the business. The objective is to allow for the dynamic mapping of identities to data sources in real time, thus minimizing any delay in making data available to authorized users.

## 4.1 Vision

Securing the On Demand Operating Environment requires a distributed security mechanism for authentication and access control within and across organizational boundaries. To provide consistent access, the security architecture for the On Demand Operating Environment supports, integrates, and unifies popular security models, mechanisms, protocols, platforms, and technologies to enable a variety of systems to interoperate securely.

### 4.1.1 Web Services security components

The security of the On Demand Operating Environment is built upon the Web Services security specifications. For more details, see the Web Services Security Roadmap at:

<http://www.ibm.com/developerworks/security/library/ws-secmmap/>

Specific bindings for security will provide protocol-specific details and security functions such as confidentiality, integrity, and authentication. For instance, transport bindings over a Secure Socket Layer (SSL) connection provide a secure connection between two end points: WS-Security is used to secure SOAP messages and WS-SecureConversation establishes a security context and derives session keys.

As businesses become more dynamic, so do their security, configuration, and user provisioning requirements. In a dynamically provisioned environment, new users may need to be added or deleted. Group access control lists may need changes. Such changes need to be reflected within the infrastructure, and synchronized, so that the security policies can be effective immediately.

Similarly, policy changes such as the granting or revocation of access rights, and changes in membership and profiles that affect access rights, should be capable of being changed dynamically, enabling runtime decisions in real time to allow pervasive yet selective and secure access to resources. Web Services-based specifications such as WS-Policy and WS-Security Policy provide the basis to unify already existing policy models.

### 4.1.2 Hardware and software security mechanisms

In addition to these capabilities, other hardware and operating system security mechanisms and facilities should also be considered.



One tremendous feature of @server systems is their secure virtualization capability. This allows server consolidation, and at the same time, secure separation of data and workloads so that even though different workloads are running on the same server, they act from a security perspective like separate boxes. pSeries and zSeries have completed Common Criteria certifications of their virtualization capabilities.

There are a few driving reasons that clients can benefit from hardware encryption in an on demand world. Encryption keys are used to protect secure communication. This may be business-to-business or consumer-to-business communication. Examples are Secure Sockets Layer (SSL) encrypting a piece of data, a Virtual Private Network (VPN) encrypting a tunnel of data, or an ATM network using encryption to connect the PIN to the identity of the customer.

Providing this type of security is very performance intensive. Hardware encryption can be used to offload performance intense cryptographic work to special processors that are designed to for this type of computations. In addition, hardware encryption can be used to provide stronger protection of encryption keys. The hardware encryption devices that are available through @server systems have been rated at the highest level of the Federal Information Processing Standard (FIPS 140). @server hardware encryption can respond to attacks and take action to prevent loss of secrets. In some industries and within some governments this level of protection is required for certain workloads.

The level of protection that is provided by the operating system is an important element in overall security. Assuring user separation, access control, authorization, and auditing of local resources are building blocks in an overall security architecture. Most of the @server operating systems have achieved third-party certifications (Common Criteria EAL certifications against CAPP or US Government Orange Book B1/C2).

To assist customers in setting up their server security, an @server Security Planner is available at:

<http://www.ibm.com/servers/security/planner/>

In addition, there are several self-protecting capabilities within the @server operating systems. For example, z/OS has built-in intrusion detection capabilities within the Communication Server to detect and defend against Denial-of-Service attacks. OS/400 and AIX sign their operating systems to assure non-authorized changes are identified. Network security capabilities of Secure Sockets Layer (SSL) and Virtual Private Network (VPN) are shipped with z/OS, AIX, and OS/400.

### 4.1.3 Glimpse of the future

Some of the directions that we see security taking in the future will enhance the On Demand Operating Environment and bring even more benefits to enterprises that are prepared to take advantage of them.

A first advance would enhance user administration with real federated identity management; users declared on one computer would have Single Sign On capability to every system they need to access.

Another advance involves inter-application security; this is where the Web Services security standards come into play. When an application tries to access a Web Service, it will have the possibility to request a security token. This token could be a Security Assertion Markup Language (SAML) authorization token issued by a trusted organization, or this could also be a more traditional certificate, based on keys.

Tivoli Access Manager v4.1 and v5.1 supports SAML and Web Services.

The IBM's Federated Identity Management (FIM) solution extends identity management for both the identity provider and service provider infrastructure. IBM Tivoli federated identity management solution builds on the current Tivoli identity and security offerings. In a federated environment, a user can log on through his identity provider in order to conduct transactions or easily access resources in external domains.

Partners in a federated identity management environment depend on each other to authenticate their respective users and vouch for their access to services. Federated identity standards, such as those being produced by the Liberty Alliance or the Web services security specifications, form an encapsulation layer over local identity and security environments of different domains. This encapsulation layer provides the ingredients for inter operability between disparate security systems inside and across domains, thus enabling federation. For more information, please refer to the redbook, *Federated Identity Management with IBM Tivoli Security Solutions*, SG24-6394.

## 4.2 How to get started today

Refer to Chapter 15, "How to secure access and control of information, resources, and applications" on page 167, for a more complete discussion of this approach.



## **How to provide scalable and consistent management and control of operations for end-to-end business systems**

System failures and poor performance can diminish employee productivity, reduce customer satisfaction, and possibly even lose customers. The key to avoiding such failures and maintaining peak performance lies in diagnosing and resolving issues quickly, before they become serious problems.

Highly interconnected systems often make this diagnosis difficult, inefficient, and expensive. A failure in a single component, such as a router, can quickly cause entire business systems to fail, generating many symptomatic alerts from multiple downstream components. The resulting confusion can cause wasted time and effort.

Business processes and operations are linked across the entire company from people all the way through the IT infrastructure. However, traditional system management tools only focus on a particular component, while the linkage between all components and the business impact of a particular component failure is not easily and quickly revealed. Enterprises need capabilities that provide end-to-end visibility for management and control operations encompassing the IT infrastructure as well as business processes. Single point management and control operations can provide cost savings and other efficiencies that make the enterprise more responsive.

## 5.1 Vision

A key part of the On Demand Operating Environment is the optimization of IT by optimizing resources and managing complexity. This will bring systems in line with business goals.

- ▶ Automating IT processes help customers meet the needs of their businesses by making the IT organizations more responsive to evolving business priorities.
- ▶ Virtualization provides for IT simplification through enterprise-wide fabrics, grids and the virtualization of internal and external resources.

These processes are implemented through the key following strategic approach.

### 5.1.1 The role of standards

Management and control technologies and standards continue to evolve. In particular, standards are being developed that will accelerate the possibility of providing a single management console and a common set of management operations. Common operations such as “stop” and “start” will be made common across the diverse set of resources in IT deployments.

Following are some examples of the IBM effort to define standards in the systems management arena, and some of their exploiters:

- ▶ IBM is actively involved in the definition of a homogeneous schema for all events, and the Common Base Event format (CBE) has been submitted to OASIS Standards body in 8/2003. This standard, developed by the IBM Autonomic Computing Initiative, supports encoding of logging, tracing, management, and business events using a common XML-based format, making it possible to correlate different types of events that originate from different applications. The common event infrastructure currently supports version 1.0.1 of the specification.

- ▶ The Integrated Solutions Console (ISC) provides a Web-based infrastructure based on industry-standard technologies (J2EE and JSR168) to address the need for common system administration in a customer's IT environment, such as setup, configuration, run-time monitoring and control, with a consistent look and feel. The solution will use WebSphere Portal and portlets as a basis. One of the exploiters of the ISC is the Virtualization Engine Console, which provides a consolidated view of enabled enterprise resources and a launch point for Virtualization Engine system services, as described in “Virtualization Engine Console” on page 57.
- ▶ Application Resource Management (ARM) is another key standard for measuring service levels and thus being able to report and to load balance in a heterogeneous infrastructure. Applications using ARM define transactions that are meaningful within the business process. Applications call ARM when transactions start and/or stop. The agent in turn communicates with management applications which provide analysis and reporting of the data. A management agent collects the status and response time.

An important capability of ARM API and the ARM agent is the tracking of hierarchical relationships among transactions. The ability to associate end-to-end transactions with units of work depends on the passing of a trace object on ARM calls. Knowledge of the parent-child relationships among transactions and the response times for each transaction enables an administrator to determine which transactions are delaying other transactions.

The ARM API is a set of standard API calls, agreed on by a number of companies. Business applications use ARM by creating objects that implement the interfaces in packages and then executing methods of the objects. From the application developer's perspective, ARM is a set of interfaces that the application loads and calls.

WebSphere Application Server, DB2 Universal Database, mySAP IDI, and Web servers, provide ARM support:

- For example, WebSphere Application Server provides ARM support for applications residing on WebSphere servers. That means it is not necessary for a programmer to instrument the applications that reside on the V5 server with ARM function calls. The administrator can configure WebSphere Application Server to make ARM calls on behalf of the applications: Each application component (JSPs, servlets, EJBs) can be wrapped with ARM tags and so be considered as a transaction for ARM measurement.
- Enterprise Workload Manager (EWLM), ARM and instrumented applications work together to provide end-to-end response times, and application/transaction topology using correlation.

The operating systems need to understand these ARM calls: all the IBM operating systems today (at the right level) include this EWLM ARM support.

### 5.1.2 Orchestration and provisioning

IBM Tivoli Intelligent ThinkDynamic Orchestrator will expand support of applications, middleware, and operating systems while at the same time providing a centralized application for enterprise-wide provisioning.

In fact, IBM servers are fully enabled for provisioning and seamlessly plug into provisioning orchestration as delivered through Tivoli Intelligent ThinkDynamic Orchestrator so that higher utilization rates can be accomplished in a short time frame.

### 5.1.3 The Service Oriented Architecture model

As specifications based on the Open Grid Services Architecture (OGSA) continue to emerge, an IT-level Service-Oriented Architecture (SOA) will become widespread. OGSA-based IT-level SOA leverages Web/Grid services to both harness internal computing resources and lease external resources on demand.

This virtualization is fundamental to the flexible delivery of IT service with availability, performance, and access to information that insulates applications and users from outages or performance bottlenecks. Dynamic reconfiguration to relocate applications, resources, and data to other elements of the distributed IT infrastructure is attainable through service abstraction and indirect interaction with underlying resources. OGSI, a core component of OGSA, is available now and provides a foundation for building Grid services and architecting SOAs.

Another fundamental value of leveraging an SOA model is to support virtualization that simplifies automation in a heterogeneous environment. It enables the use of a common set of services that in turn utilize internal adapters to access legacy resources.

Virtualization through traditional hypervisor technologies provides the valuable function of projecting multiple time-shared logical images onto the same physical resource. This form of virtualization is completely complementary to the distributed virtualization capabilities of an applied SOA.

Managing abstract, distributed, and heterogeneous resources as one big common system is exactly the kind of problem that Grid technology can solve. Just as an Internet user views a unified instance of content via the World Wide Web, a Grid user essentially sees a single, large, virtual computer.

## 5.1.4 Automation

The following aspects, in regard to automation capabilities, need to be considered.

### **Workload management**

Workload management will evolve from the current zSeries WLM capabilities and Enterprise Workload Manager (EWLM) capabilities, by integrating existing and future management products with enhancements for more sophisticated levels of monitoring, managing and reporting. EWLM provides the ability to monitor a transaction across heterogeneous servers and to correlate the events, resulting in a quicker and more accurate understanding of problems affecting groups of transactions. In the future, EWLM will help to manage more resources

The management and control operations at every step and every layer in a business process will become policy-based and seamlessly span resources from people to IT infrastructure components.

### **Tivoli System Automation**

The Tivoli System Automation product family, which is used to automate and manage the availability of business applications in homogenous clusters, will move towards a policy-based autonomic end-to-end automation for e-business applications spanning multiple platforms. This will allow for not only automation of single tiers of an e-business application, but also its management end-to-end.

Through common event infrastructures, both local and solution-level problem determination will become easier and in turn enable more autonomic management of resources. These common event infrastructures and symptom formats also enable self-healing services such as call-home.

### **Grid services**

Grid services provide for the controlled management of the distributed and often long-lived state that is commonly required in sophisticated distributed applications. OGSA also introduces standard factory and registration interfaces for creating and discovering grid services. There will be continued evolution of grid and autonomic functions facilitating the integration and management of server, OS, application, network and storage security, optimization, provisioning, billing, metering, and so on. Grid services broaden the scope of managed resources beyond traditional machines. Managed resources include much finer grained resources such as network devices and storage devices. OGSA will provide a *lingua franca* for managing a wide variety of both logical and physical resources.

## 5.2 How to get started today

Important aspects of the control and management of end-to-end systems are the provisioning of operating systems, applications, and other resources; monitoring of the availability of each component; and the ability to control and manage operations at every step and every layer in the business process across disparate applications and IT infrastructures.

These components are demonstrated in the On Demand Operating Environment framework as Policy-Based Orchestration, Availability, and Provisioning. Orchestration refers to the ability to:

- ▶ Instantiate and enforce business policies for automated change
- ▶ Coordinate across core automation disciplines
- ▶ Automate allocation of resources to priority applications
- ▶ Assure application service levels under peak demands
- ▶ Configure levels of automation for evolutionary adoption

Business systems typically span Web-based, client-server, and host environments; are comprised of many interconnected application components; and rely on diverse middleware, databases, and supporting platforms.

There are several choices for customers who need to achieve high level control and management of end-to-end systems for availability, orchestration and provisioning. In order to achieve application performance goals, automation is key. The following sections depict some of the key components to start with in this area.

### 5.2.1 Enterprise Workload Management component

The EWLM end user interface, EWLM Control Center, provides a Web based interface which provides a consolidated view of the transactions and their servers that they run on, regardless of the operating environment. Therefore, EWLM presents the CPU of each server and middleware product and the application topology for the transaction as it traverses from component to component (including the time the transaction spent in each of the middleware components or waiting on a component) across the heterogeneous environment. As additional systems are dynamically added (or taken away), EWLM will discover the new resource in its topology view and immediately begin leveraging it to rebalance the workload in order to achieve the performance objective.



Although today EWLM is limited to monitoring business objectives, it will evolve into a real-time manager of system resources (CPU, dispatching priority, etc.). Based on the service class goals in the EWLM policy, EWLM will learn about the workloads running, determine if a change is necessary when the goals are not being met and determine if the changes will not only have a positive affect on the workload, but determine what impact the change will have on other workloads that the resource is being taken away from. Based on workload prioritization, EWLM will modify the resources needed. Then, once the modification has been made, EWLM will continue to monitor the workloads and adjust the resources necessary in order to provide an automated management function for end-to-end business systems.

### **5.2.2 IBM Tivoli Business Systems Manager component**

IBM Tivoli Business Systems Manager provides a single point of management and control for real-time operations for end-to-end business systems management. It enables customers to graphically monitor and control interconnected business components and operating system resources from one single console and give business context to management information and decisions. It helps manage business systems by understanding and managing the dependencies between business system components and the underlying infrastructure.

### **5.2.3 Virtualization Engine Console component**

The Web based integrated solutions console which comes with the IBM Virtualization Engine Suite for Servers is called the Virtualization Engine Console. Besides providing a launchpad for other Web based consoles, including other Virtualization Engine system services, such as EWLM, the Virtualization Engine Console provides system administrators with a consolidated view of their environment which they are responsible for.

The health center provided with the Virtualization Engine Console provides administrators with the ability to monitor and manage servers across multiple platforms (AIX, Linux, OS/400, and Windows) from a single interface. The health center simplifies your work as a system administrator by providing a Web based console that connects to and displays the health of your cross-platform @server System and Storage management applications, such as Cluster System Management, IBM Director Multiplatform and IBM Director, Management Central, and IBM Tivoli Monitoring. You can identify problems across your system and storage environment and take actions to solve these problems from the health center by simply clicking on the resource, monitor, or task that needs attention.

## 5.2.4 IBM Director Multiplatform component

Often, administrators manage AIX, Linux, and Windows servers as individual physical servers. By consolidating to IBM @server BladeCenter technology, companies can realize immediate benefits; however, the problem of managing the respective AIX, Linux, and Windows servers and populating new blades still exists. IBM Director Multiplatform unites those disparate operating systems into a single, user-friendly view.

It automates such systems management tasks as resource monitoring, task health, corrective management, and console launching. Its inventory functions discover all the hardware available in the environment, its configuration, and its status by scanning each managed system. Automated problem determination capabilities point out problems with hardware, such as faulty power supplies, fans, voltage regulator modules, and network interface cards, among others.

Administrators can monitor specific processes and system resources and receive alerts when stated thresholds are reached or when a process does not start. IBM Director Multiplatform allows administrators to define event triggers separately from actions, providing the capability to create custom action plans for various systems and groups of managed systems. Its mass configuration capability can configure SNMP community names and trap destinations for multiple systems at once.

IBM Director Multiplatform can easily deploy asset identification across machines and remotely perform network operations such as specifying domains, adding DNS servers to network properties, and setting DHCP on managed systems. Through its remote management, automated problem determination, and self-monitoring autonomic computing capabilities, IBM Director Multiplatform increases the availability of managed servers as well as the efficiency of administrators by reducing the down time of hardware and managed systems.

## 5.2.5 IBM Tivoli Enterprise Console component

IBM Tivoli Enterprise Console® helps enterprises move from monitoring and control to problem diagnosis and resolution, thereby improving system performance and reducing support costs. It provides an auto-discovery feature to identify the components that make up the environment. It then correlates and processes events based on pre-configured rules that provide best-practices event management out-of-the-box. The integration with IBM Tivoli Risk Manager provides monitoring and management of firewalls and intrusion detection systems, and IBM Tivoli Comprehensive Network Address Translator enables integrated management of overlapping IP domains.

## **5.2.6 IBM Tivoli Monitoring components**

IBM Tivoli Monitoring provides monitoring for essential system resources, to detect bottlenecks and potential problems, and to automatically recover from critical situations. Tivoli Monitoring saves system administrators from manually scanning through extensive performance data before problems can be solved. Using industry best-practices, Tivoli Monitoring can provide immediate value to the enterprise. This offering spans operating systems; pre-built packs are available to support application servers, databases, other middleware, and packaged applications.

## **5.2.7 IBM Tivoli Intelligent ThinkDynamic Orchestrator component**

IBM Tivoli Intelligent ThinkDynamic Orchestrator provides dynamic policy-based provisioning of systems, middleware, and applications. It can gather information about the performance of application clusters and build a workload model that can predict impending resource requirements; manage resources across application clusters to optimize business-aligned service delivery; automate the deployment of computing resources to each application environment; and provide applications with priority access to data center resources based on class of service.

Business applications are automated and made highly available using the IBM Tivoli System Automation product family.

## **5.2.8 The storage components**

Today's storage management requirements go beyond traditional backup and recovery solutions. Data is the currency of today's e-business economy, and a storage solution must encompass data reliability, solution scalability, and disaster planning and recovery. The storage management solution impacts the overall infrastructure as well as individual mission-critical applications.

Heterogeneous storage infrastructures, driven by growth in files and database data, are consuming increasing amounts of administrative time. IT managers are looking for ways to make their storage administrators more efficient. The IBM implementation of Storage Infrastructure Management brings together four "subject matter experts," each designed to perform complete management in its area of expertise, and empowers administrators with automated tools for managing heterogeneous storage infrastructures.

These four products, which are packaged under IBM TotalStorage Productivity Center, are described in the following sections.

### **IBM TotalStorage Productivity Center for Fabric**

IBM TotalStorage Productivity Center for Fabric (formerly IBM Tivoli Storage Area Network Manager) is the expert on the SAN fabric that connects the host systems and applications to the storage devices to create, by correlation between different sources of information, topology mapping of the SANs that can be displayed graphically from a central management point.

### **IBM TotalStorage Productivity Center for Data**

IBM TotalStorage Productivity Center for Data (formerly IBM Tivoli Storage Resource Manager) is the subject matter expert on the data that is stored in the infrastructure: the files, file systems, databases, and tablespaces. It uses a global Policy-based management that enable you to define and enforce storage policies through user-defined alerts, quotas, and constraints, notifying the user via e-mail, pager, or the event log, or a systems management console for events such as when a quota has been exceeded or a constraint violated. IBM TotalStorage Productivity Center for Data can use also the Event Integration Facility (EIF) to send messages to the Tivoli Enterprise Console (TEC). This can allow TEC to consider IBM TotalStorage Productivity Center for Data alerts in causal analysis for problems. TEC will be added as a destination for alerts, in addition to SNMP Trap, Windows Event Log, and so forth.

### **IBM TotalStorage Productivity Center for Disk**

IBM TotalStorage Productivity Center for Disk (formerly IBM TotalStorage Multiple Device Manager — Performance Manager feature) is a centralized management console of networked storage devices such as ESS and DS4000 product families or also SAN Volume Controller.

### **IBM TotalStorage Productivity Center for Replication**

IBM TotalStorage Productivity Center for Replication (formerly IBM TotalStorage Multiple Device Manager — Replication Manager feature) is centralized management console of replication on network storage devices such as the IBM TotalStorage Enterprise Storage Server 'ESS.

## **5.2.9 The Electronic Service Agent component**

The Electronic Service Agent™ (ESA) is a tool that resides on a system to monitor events and transmit system inventory information to IBM on a periodic, customer-definable timetable. This monitor tracks and captures system inventory, hardware error logs, and performance information, automatically reporting hardware problems to IBM. Information collected through this ESA is available to IBM service support to assist in diagnosing problems. With early knowledge about potential problems, IBM can proactively respond to customers and assist in maintaining higher availability and performance.

## 5.2.10 The automation environment

The following components need to be considered for their automation capabilities.

### **IBM Tivoli System Automation**

Business applications that are automated by Tivoli System Automation can run in a single node environment or in a clustered environment. As an example, IBM Tivoli System Automation can be used to automate a database such as DB2 or a full complex SAP installation.

The IBM Tivoli System Automation product family covers the broad spectrum of IBM's @server platforms. For the distributed world, IBM Tivoli System Automation is available on Linux for xSeries, pSeries, iSeries, and zSeries.

IBM Tivoli System Automation is a family of products that manages the availability of business applications by performing automated, goal-based problem resolution in failure situations. By monitoring business applications and their related IT resources, IBM Tivoli System Automation products can detect outages quickly and react with automated actions for rapid and consistent recovery of failed resources and whole applications.

### **Autonomic Computing Toolkit**

The Autonomic Computing Toolkit is a set of components that simplifies the incorporation of autonomic functions into applications. The technologies provided in the Autonomic Computing Toolkit can be used to develop or enhance certain capabilities in products and systems to facilitate end-to-end system management. These capabilities include problem determination, common systems administration, and solution installation and deployment.

Problem determination autonomic capabilities can be developed with the Autonomic Management Engine (AME), the Generic Log Adapter, and the Log and Trace Analyzer tool. The Integrated Solutions Console is used to build effective common systems administration capabilities. The dependency checker is a technology that provides autonomic capabilities for solution installation and deployment. The Autonomic Computing Toolkit is available at:

<http://www-106.ibm.com/developerworksautonomic/overview.html>

## Workload management

Multiple workload management capabilities exist today:

- ▶ The Workload Manager for z/OS and S/390 component allows for the definition of performance goals for a workload and assigns a business importance to each goal in a zSeries environment. Through WLM, goals are defined in business terms, and the system decides how much resource, such as CPU and storage, should be given to it to meet the goal. Workload Manager constantly monitors the system and adapts processing to meet the defined goals. WLM's scope extends from the network, with intelligent TCP/IP and SNA routing, within a single z/OS image, out to the disks where the data is stored. WLM also manages priorities across multiple systems in a Parallel Sysplex.

Once the goals have been defined to WLM, the RMF™ component of z/OS works in concert with WLM to report goal achievement and identify the causes of delays by workload as well as by component. This not only enables the reporting of SLA attainment, it also provides the ability to identify the causes for specific applications missing their goals on z/OS systems.

- ▶ IBM Enterprise Workload Management (EWLM), a component of IBM Virtualization Engine Suite for Servers, extends the scope of workload management capabilities across platforms, to monitor applications and resources end-to-end. EWLM is planned to deliver workload monitoring capability to enable the shifting of network resources and workloads based on changing requirements. EWLM is a breakthroughs technology that brings self-optimization, in the form of performance and response time management, to an entire Virtualization Engine-enabled IT environment.

WLM and z/OS WLM can run simultaneously; EWLM monitoring does not affect z/OS management and monitoring. A z/OS customer can continue z/OS management and use EWLM to see how the work relates to transactions that exist within a network of different platforms, such as AIX, OS/400, Linux, or Windows.

- ▶ With the pSeries, starting with AIX V5.2, the AIX Partition Load Manager (PLM) is a tool providing automated CPU and memory resource management among AIX partitions running on Power5 systems and optimizes CPU and memory resource use by AIX partitions; PLM monitors and reconfigures CPU and memory resources across a group of AIX LPARs; it uses defined policies which specify how resources are allocated. PLM is a component of the Power5 Advanced Power Virtualization option.


## 5.3 Key products to start with

The following IBM products should be evaluated for their use in developing a solution with this approach:

- ▶ IBM Tivoli Access Manager for Business Integration (prior to March 2003, this product was known as IBM Tivoli Business Systems Manager)
- ▶ IBM Tivoli Risk Manager
- ▶ IBM Tivoli Enterprise Console
- ▶ IBM Tivoli Monitoring
- ▶ IBM Tivoli Comprehensive Network Address Translator
- ▶ IBM Tivoli Intelligent ThinkDynamic Orchestrator
- ▶ IBM Tivoli System Automation product family
- ▶ IBM Tivoli Decision Support
- ▶ IBM Report Management Facility (RMF) component of z/OS
- ▶ IBM Workload Manager for z/OS
- ▶ IBM Virtualization Engine Suite for Servers, including:
  - IBM Virtualization Engine Console
  - IBM Enterprise Workload Manager
  - IBM Director Multiplatform
  - IBM Grid Toolbox
- ▶ IBM Electronic Service Agent
- ▶ IBM Tivoli Storage Manager
- ▶ IBM Tivoli Storage Resource Manager
- ▶ IBM Tivoli Storage Area Network Manager
- ▶ IBM Multiple Device Manager
- ▶ IBM Autonomic Computing Toolkit
- ▶ IBM TotalStorage Productivity Center products







## How to avoid system failures and take automated action to resolve problems

The IT environment relies on a large number of systems with an increasing amount of resources, such as processors, memory banks, hard drives, cables, switches, and so on. Even if the progress in technology allows all those components to be more and more reliable, failures will still occur.

The On Demand Operating Environment needs to anticipate problems and enable automated policy-based approaches to problem resolution. The overall objective and purpose of availability management is to provide a policy-based declaration of level of resiliency for a deployed solution.

The capabilities of the On Demand Operating Environment will be exploited to predict and monitor failures in the system, and to preempt, fix, react, and respond to those failures and degradations. Actions may range from a simple fix for a failure, to more complex schemes for redistribution of workload away from failed systems, or even extreme actions that move or re-deploy the entire underlying infrastructure to meet availability requirements. Tooling and services are required to understand the resource composition of an operating environment in order to re-provision those resources.

Another important aspect in the avoidance of system failures is through virtualization. By masking (or hiding) the physical resources from the end user and providing a virtual interface, certain system, network, storage and other failures can be avoided by the end user while automation actions can be implemented to flag and/or fix the failed area. These automation actions would include such things as automatic failover, provisioning, and problem avoidance.

## **6.1 Vision**

Web services performing automated problem determination and corrective actions will be available through the Enterprise Service Bus. They will be in the form of middleware that supports pluggable policies and rules that determine the behavior of business processes across all tiers of its implementation. For example, a simple policy can state that a particular WebSphere application can be readily re-started if needed to optimize an overall process.

These capabilities will be enabled by using common event infrastructures and symptom formats. Autonomic technologies will play a key role in avoiding system failures and automating corrective actions.

The combination of dynamic provisioning and automation will allow the integration of automated problem resolution with dynamic and optimized provisioning to help ensure the meeting of SLAs. A practical example has to do with handling hardware failures or disasters. In such cases, software-level “healing” is clearly not an option and internal or external provisioning is required.

A continued focus on business resiliency policies will allow integrated billing, metering, and availability solutions. This coming together of business process information combined with declaratively described policies for problem resolution will allow for optimizing the attainment of service level objectives.

## **6.2 How to get started today**

A complete solution requires the integration of elements from hardware, software, and services. Of course, these may be incrementally deployed based on the immediate needs of the enterprise. The following sections describe the main solutions that can be implemented.

### **6.2.1 Hardware solutions**

In the following paragraphs we discuss some of the hardware solutions that can be deployed today to avoid system failures.

## **xSeries environment**

From the xSeries brand, we pick the active memory, included in high-end servers. This feature along with the correct level of operating system allows the administrator to enable memory mirroring. This allows all write operations to be transparently made to two different memory banks. The event of a failure would go unnoticed by the operating system and it is even possible to hot swap the defective component. IBM also provides the chipkill feature, which allows the memory to be protected against chip failure, with RAID-5-like capabilities and spare memory chips. This chipkill feature is also available on the other hardware brands and makes the memory subsystem about 1000 times more reliable than the usual ECC protection. Other virtualization technologies include an integrated shared infrastructure for Blades, VMWare, clustering, Virtual Machine Manager, Capacity Manager, and a SANFS Client.

## **iSeries environment**

The iSeries overall hardware is designed so most of the components are redundant and hot swappable. For example, power supplies are hot swappable with batteries providing enough power to operate the system while the power supply is replaced. PCI slots allow a failing adapter to be replaced without having to stop applications. The iSeries also has virtualization technologies such as dynamic LPAR capability, clustering, virtual Ethernet, virtual I/O, micro-partitioning, capacity on demand, IBM Director, and VE Console support.

## **pSeries environment**

An example from pSeries is processor dynamic de-allocation. During normal operation the processors are constantly monitored and any retry is noticed and logged. After a given threshold has been reached, the decision is automatically taken that the given processor is no longer reliable and should be deactivated. The removal of the processor is done automatically, without any human intervention, and without stopping the operating system. Using the latest version of the AIX operating system, and if any Capacity upgrade On Demand (CUoD) processors are available on the system, one of the inactive processors will then be picked to replace the failing one, and the system can return to normal operation. Other pSeries examples are clustering, NIM, Virtual Ethernet and Virtual I/O support, SANFS Client, and IBM Director.

## **zSeries environment**

In the zSeries environment, all CPCs are delivered with spare processors that can be used to transparently replace a failing one. The zSeries range also supports Capacity Upgrade on Demand (allowing non-disruptive upgrades), Capacity BackUp (allowing transparent upgrades to be completed within minutes) and Daily On/Off (allowing additional processors to be turned on and off as required).

In addition, Parallel Sysplex technology, where up to 32 z/OS systems behave like a single, logical computing facility, can be exploited. The underlying structure of the Parallel Sysplex remains virtually transparent to users, networks, applications, and even operations. The HyperSwap™ support delivered by the Geographically Dispersed Parallel Sysplex (GDPS) offering makes it possible for application availability to be maintained even across a complete planned or unplanned site outage, potentially masking the failure to application users. There is also a zSeries Application Assist Processor (zAAP), hipersockets support, Virtual LANs, and z/VM Virtual Machine Resource Manager, Intelligent Resource Director (IRD), Workload Manager (WLM) and IBM Director support

Across all of the IBM @server brands, there is also the ability to host a “call home” function. When a failure occurs, it can be automatically reported to IBM maintenance services, which automatically orders the replacement part and sends someone to replace the failing component.

## 6.2.2 Software and services solutions

The following paragraphs describe some of the software and services solutions that can be deployed today to avoid system failures.

### **Task management**

The following component must be considered for task management.

#### ***IBM Director Multiplatform***

IBM Director Multiplatform helps deliver a common, consistent, cross-platform systems management solution for IBM servers, storage, and operating systems. It provides a single administrative console for management tasks (operating system, storage management, distributed systems management, and platform management), a common management infrastructure for upward integration with Tivoli\*, and a management foundation for the on demand architecture.

Using IBM Director Multiplatform, many of the administrators manual tasks can be automated to proactively and remotely manage systems. Tasks such as discovery, event logs and action plans, file transfer, inventory collection, process management, resource monitors and thresholds are among those tasks that fall into this category. Additionally, the predictive and proactive capabilities associated with alerting and real-time system diagnostics help maximize server uptime and reduce service downtime costs.

IBM Director Multiplatform also provides a bridge to the Virtualization Engine console. This bridge enables uses of the Virtualization Engine console health center to monitor and take some corrective actions from a task-oriented, Web-based interface.

## **Performance goal monitoring**

The following components must be considered to monitor performance according to goals.

### ***Enterprise Workload Management***

Enterprise Workload Manager (EWLM) provides the ability to monitor the performance goals that you have defined for your workload based on your business policy or Service Level Agreements. It provides topology views for the application and servers for which the transaction traverses and provides information that may highlight the potential problem area (for example, a communication problem between WebSphere Application Server and DB2).

Additionally, EWLM monitors the health and load of the servers and their applications and makes decisions as to which server or application is the best to handle new requests. Currently, this information can be used to send recommendations to load balancers, such as Cisco's load balancing product to influence future routing. See "How to balance workloads in the network" on page 255 for a further description on how EWLM and Cisco provide effective load balancing solutions which automatically avoids poor performing servers, thus avoiding potential problems. In the future, EWLM will expand its capability from monitoring to actively managing resources for resource optimization and problem avoidance.

EWLM allows you to define transactions to be associated with an EWLM policy service class, and therefore EWLM goals for each service class. EWLM looks at the end-to-end transaction and determines whether or not the transaction goals are being met. An IT administrator can use the EWLM views to observe potential performance problems for critical applications, and by using the detailed performance views, determination may be made as to the root cause of the problem.

Alternatively, based on the information provided by EWLM, the IT architect can then use *Tivoli Monitoring for Transaction Performance* (TMTP) to determine the root cause of the problem. The IT architect can then use *Tivoli Provisioning Manager* (TPM) to execute a workflow to move a free resource from a pool into the constrained environment to resolve the performance problem. EWLM discovers the new resource in its topology view and immediately begins leveraging it to rebalance the workload in order to achieve the performance objective.

### ***IBM Tivoli Monitoring***

IBM Tivoli Monitoring provides monitoring for essential system resources, to detect bottlenecks and potential problems, and to automatically recover from critical situations. Tivoli Monitoring saves system administrators from manually scanning through extensive performance data before problems can be identified

and solved. Using industry best-practices, Tivoli Monitoring can provide immediate value to the enterprise.

IBM Tivoli Monitoring also provides the foundation for additional automated best practices via Proactive Analysis Components (PACs) for managing business-critical hardware and software including middleware, applications, and databases. IBM Tivoli Monitoring provides seamless integration with other solutions. Combined with IBM Tivoli Business System Manager and IBM Tivoli Enterprise Console, it provides a true end-to-end solution.

## **Automation**

The following components need to be considered for automation.

### ***IBM Tivoli System Automation***

IBM Tivoli System Automation is a family of products that manages the availability of business applications by performing automated, goal-based problem resolution in failure situations. By monitoring business applications and their related IT resources, IBM Tivoli System Automation products can detect outages fast, and react with automated actions for a quick and consistent recovery of failed resources and whole applications.

Business applications which are automated by Tivoli System Automation can run in a single node environment or in a clustered environment. As an example, IBM Tivoli System Automation can be used to automate a database such as DB2 or a full complex SAP installation.

The IBM Tivoli System Automation product family covers the broad spectrum of IBM's @server platforms. For the distributed world, IBM Tivoli System Automation is available on Linux for xSeries, pSeries, iSeries, and zSeries. IBM Tivoli System Automation for Linux brings out-of-the-box support for a large number of applications by its pre-configured high availability policies.

For the mainframe environment, *IBM Tivoli System Automation for OS/390* can help customers with single z/OS or OS/390 systems and Parallel Sysplex clusters to simplify management, reduce costs, and increase availability. Tivoli System Automation for OS/390 is designed to automate I/O, processor, and system operations and includes “canned” automation for IMS™, CICS, IBM Tivoli Workload Scheduler, and DB2. As a potential next step, the IBM Tivoli System Automation product family will move towards autonomic end-to-end automation for e-business applications by integrating heterogeneous platforms under one umbrella.

### ***IBM Tivoli Intelligent ThinkDynamic Orchestrator***

IBM Tivoli Intelligent ThinkDynamic Orchestrator reduces costs and improves server utilization. It helps boost server-to-administrator ratios by automating the steps to provision (or re-provision in the case of failure recovery), configure, and deploy a solution into production. This automated process supports servers, operating systems, middleware, applications, and network devices acting as firewalls, routers, switches, and load balancers. By utilizing existing hardware, software, and network devices without rewiring, implementation times are minimized and a faster return on investment is achieved.

### **Storage environment**

The following components must be considered for the storage infrastructure environment.

#### ***IBM TotalStorage Productivity Center for Fabric***

IBM TotalStorage Productivity Center for Fabric (formerly IBM Tivoli Storage Area Network Manager) provides a feature called the SAN Error Predictor that monitors the links in the fabric and predicts the most likely hardware failures.

This ED/FI (Error Detection and Fault isolation) function is accomplished by having the Fibre Channel adapters gather error data at the hardware level. The data includes the number of data errors and the underlying causes such as loss of light, loss of synchronization, or operational readiness. This data is put into a historical database and an Analysis Engine is run against it to spot trends and do predictive failure analysis.

When the combination of errors indicates that the problem is the result of a condition that will get worse and is likely to result in a hard failure (as opposed to transient errors that can be tolerated) a notification is sent to the administrator. This allows a scheduled replacement of the failing network hardware component instead of a reactive crisis. Most importantly, in an era where acceptable levels of unavailability are measured in minutes over the course of a year, application availability can be proactively pursued.

#### ***IBM TotalStorage Productivity Center for Data***

IBM TotalStorage Productivity Center for Data (formerly IBM Tivoli Storage Resource Manager) implements an automated file system extension capability that allows the user to specify a policy to automatically extend a file system when a threshold has been reached. For example, if a file system's threshold is set at 78% and, through monitoring, IBM TotalStorage Productivity Center for Data identifies that this threshold has been exceeded, it can automatically initiate a file system extension to reduce the possibility of a storage-related outage.

## The Electronic Service Agent

The first step in deploying the solution would be to start using the electronic service agent on all the servers in the IT environment. Any server under warranty or IBM maintenance is automatically enlisted in the electronic service agent program. It is possible to have an electronic service agent gateway that will collect the alerts from servers on the network and will be the point from which all the calls to IBM maintenance will be made. It is also possible to have a backup gateway in case the first one is inoperative. Figure 6-1 shows the basic architecture deployed for the electronic service agent.

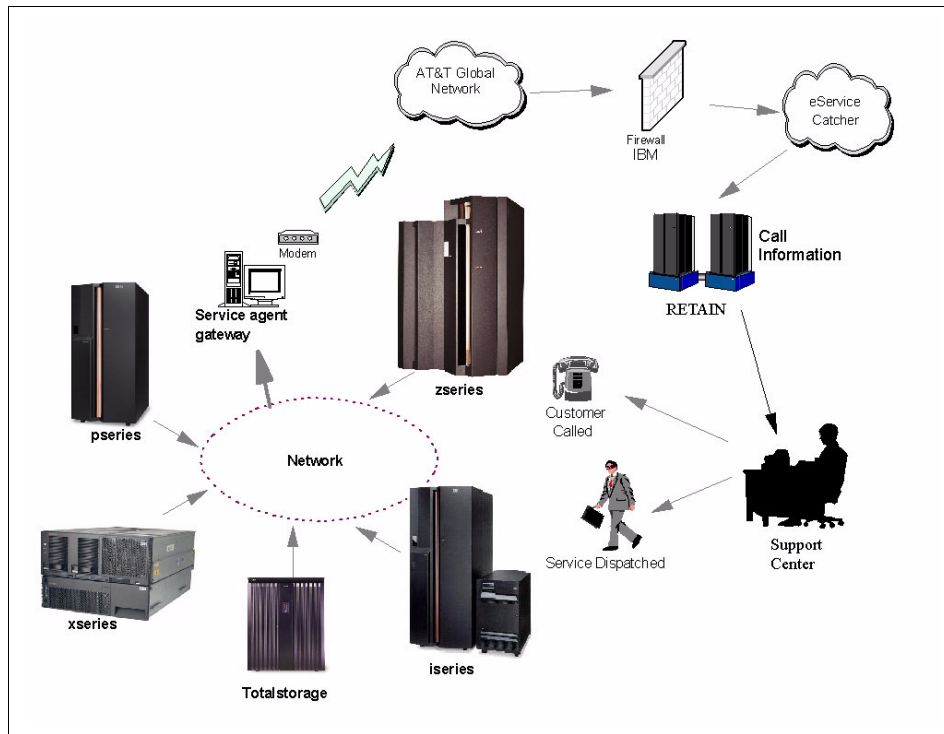


Figure 6-1 Electronic service agent architecture

Using the electronic service agent helps ensure quick recovery of hardware failures. However, that is just a first step. To centrally monitor systems to be aware of other failures (such as software), and know that service levels are being met, requires products such as IBM Tivoli Monitoring. Figure 6-2 shows the basic architecture of the solution. The first step is to install the agent (endpoint) on the servers to be monitored. The supported platforms include AIX, HP-UX, Linux, OS/400, Solaris, Windows NT®, and Windows 2000.



## IBM Tivoli Monitoring

Out of the box, IBM Tivoli Monitoring provides automated best practices, which eliminates the need to perform extensive research and manual configuration of the solution for a quick start. In a second step, the solution can be customized to support specific applications using the workbench feature. The health console enables users to view both real-time and recent historical data for any server via a Web browser or Java-based graphical user interface.

Alerts can be forwarded to the Tivoli Enterprise Console in order to centralize the monitoring of all events. Finally, the data warehouse feature provides the ability to leverage advanced technology and extend monitoring capabilities and historical reporting across the Tivoli Proactive Analysis Components for applications, databases, and middleware, such as mySAP.com, Siebel, IBM WebSphere MQ, WebSphere MQ Integrator, WebSphere MQ Workflow, WebSphere Interchange Server, DB2, Oracle, and Informix®.

Figure 6-2 shows the different components of the Tivoli Monitoring architecture.

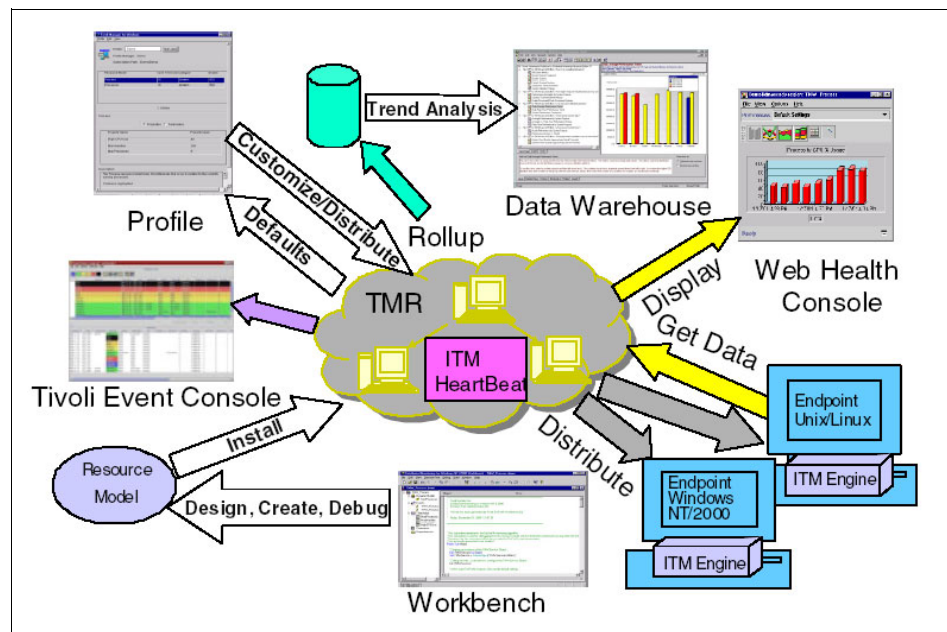


Figure 6-2 IBM Tivoli Monitoring high-level architecture

## 6.3 Key products to start with

The following IBM products should be evaluated for their use in developing a solution with this approach:

- ▶ IBM TotalStorage Productivity Center for Data
- ▶ IBM TotalStorage Productivity Center for Fabric
- ▶ IBM Tivoli Monitoring
- ▶ IBM Tivoli Business Systems Manager
- ▶ IBM Tivoli Enterprise Console
- ▶ IBM Tivoli System Automation for Linux
- ▶ IBM Tivoli Systems Automation for OS/390
- ▶ IBM Tivoli Intelligent ThinkDynamic Orchestrator
- ▶ Geographically Dispersed Parallel Sysplex
- ▶ IBM Tivoli Storage Area Network Manager
- ▶ IBM Virtualization Engine, Suite for Servers components:
  - IBM Director Multiplatform
  - Enterprise Workload Manager
  - Virtualization Engine Console



## How to protect systems from intrusions and threats using monitor and alert systems

In today's complex network environment, e-business enables significant transactions to take place online, both business-to-business (B2B) and business-to-consumers (B2C), with varying levels of trust and knowledge about each other. Business face increasing risks from a multitude of fronts such as virus threats, unauthorized access, denial of service attacks, and so on, that target networks, servers, and desktops. The risks increase as more enterprise systems and applications become accessible on the Internet. e-business has created an environment in which a technical attack can translate directly into legal liability if confidential information is exposed; a drop of customer confidence due to a downed site or exposed information; loss of income; and even a loss of brand equity. The most important concern is how to let the good guys in and keep the bad guys out.

Companies have made significant investments in many key products such as firewalls, intrusion detection systems, and application-level security, mainly because each key product implements a specific security function that is required to implement the overall security strategy. However, security is more than a firewall or other product solution. An integrated security management solution enables corporations to make the most informed security decisions by leveraging

the intelligence of the various security links. Centrally correlating intrusions and vulnerabilities across different components provides the overall assurance that individual security components reinforce and complement each other and implement the overall business goal of managing risks to information assets.

## 7.1 Vision

The On Demand Operating Environment autonomic capabilities will continue to expand. In particular, *Intrusion Detection* capabilities will expand from networking and network-based devices to additional layers including middleware and applications. For example, increased Web commerce during the holidays may be the reason for increased network traffic. Close connection between business policies, middleware, and network-level intrusion detection is required to make correlations that can, for instance, differentiate between a denial of service attack and a sharp increase in business demand.

Operating system level intrusion detection mechanisms will continue to be improved and made extensible. Integration of these technologies with higher level correlation mechanisms such as IBM Tivoli Risk Manager will continue to evolve.

There will be a new class of intelligent management for intrusion detection that is strongly tied to provisioning and optimization. For example, it may provide the ability to automatically provision a new server to deal with a denial of service attack while other servers continue to honor authentic transactions.

## 7.2 How to get started today

To address the challenges of building an on demand IT infrastructure, IBM delivers automation capabilities for On Demand Operating Environments. As part of the autonomic characteristics of the On Demand Operating Environment, the security infrastructure needs to provide the ability to achieve self-protection. This includes the ability to detect intrusions as well as react dynamically to such intrusions to prevent further damage or continued degradation of system performance. The following sections describe the main components which address these challenges.

### 7.2.1 The IBM Integrated Security Solution for Cisco Networks

In early 2004, IBM and Cisco announced the extension of an existing strategic alliance to address the end-to-end security concerns of e-businesses. The IBM and Cisco collaboration begins with a focus on four areas:

- ▶ Integrating IBM Tivoli Identity Manager and Cisco Access Control Server to help lower user life cycle management costs.
- ▶ Enhancing the security of IBM platforms by integrating Cisco endpoint security technologies with the IBM embedded security chip, a ThinkVantage Technology in Lenovo ThinkPad laptops and ThinkCentre desktops.
- ▶ Securing and connecting endpoints with systems and applications through IBM participation in the Cisco Network Admission Control (NAC) program.
- ▶ Offering customers the option of ordering the Cisco Security Agent technology with selected IBM *@server* xSeries models.

The IBM Integrated Security Solution for Cisco Networks, which became generally available at the end of 2004, identifies specific IBM products and services offerings that contribute to address some of the security concerns in a network environment. The IBM Integrated Security Solution for Cisco Networks provides automated user access to applications, network resources and data; at the same time it protects these resources from unauthorized use and non-secure devices. IBM and Cisco Systems provide products that help enterprises address security exposures in two key areas:

- ▶ The first is managing the identities of users who connect to the enterprise.
- ▶ The second is dealing with “wellness” of a device connecting to the network.

The following IBM products are involved in this solution:

- ▶ IBM Tivoli Security Compliance Manager V5.1.2, that now supports Secure Access Control Server<sup>1</sup> by providing the capability to define a compliance policy for access to the network, to execute that policy on the end system, and to communicate with the Cisco network about the compliance or lack of compliance of the end system. If the end system is not compliant, remediation can be provided by Tivoli Provisioning Manager when used in conjunction with Tivoli Security Compliance Manager
- ▶ IBM Tivoli Provisioning Manager 2.1, that offers new workflows to remediate end systems when found out of compliance by the network. These workflows provide capability for Windows patches, zone alarm firewall, password strength, Norton AntiVirus, and service packs. Additional workflows can be easily developed to encompass other remediation needs based on local policy
- ▶ IBM Tivoli Identity Manager 4.5 now supports Secure Access Control Server to provide automation of user provisioning for port level authentication. It already provides capability to define users for wireless, VPNs, and other network user requirements

---

<sup>1</sup> The Cisco Secure Access Control Server extends access security by combining authentication, user or administrator access, and policy control from a centralized networking identity solution,

- ▶ Lenovo ThinkPad and ThinkCentre Systems with Antidote Delivery Manager software
- ▶ IBM Tivoli Access Manager

IBM Integrated Security Solution for Cisco Networks is designed for clients who use Cisco routers, networks and firewalls, as well as Cisco Secure Access Control 3.3 Server, Cisco Trust Agent 1.0, Cisco IOS and IOS security image and Cisco Security Agent 4.0 (please work with an authorized Cisco representative or business partner to determine the relevant Cisco components).

### **The Cisco Network Admission Control component**

The Cisco Network Admission Control (NAC) program allows businesses to examine the security status of endpoints such as PCs and servers, and automatically permit or deny endpoint access to critical network and system resources. Admission is based on customers' pre-defined corporate IT security policies. IBM's participation in the Cisco NAC program extends the capability to automatically examine system and application credentials and provide more effective remediation strategies for non-compliant endpoints that could otherwise pose a security threat to businesses.

### **IBM Tivoli Security Compliance Manager component**

IBM Tivoli Security Compliance Manager comes with pre-defined, recommended security policies that can be customized to fit specific corporate, industry, or regulatory security policies. It acts as an early warning system by identifying security vulnerabilities and security policy violations. It collects and audits data about security status, and allows administrators to produce detailed reports. It automates security scans of servers and desktop systems, both before and during end point connections to the network.

### **IBM Tivoli Provisioning Manager component**

IBM Tivoli Provisioning Manager automates manual tasks of provisioning and configuring servers, operating systems, middleware, applications, storage, and network devices. Using workflows, it automates the remediation of noncompliant end points by installing required software updates or by correcting configuration issues. The remediation capabilities of Tivoli Provisioning Manager include software levels — typically operating system levels and fix packs; patch levels; virus protection, and firewalls. Through workflows and automation, Tivoli Provisioning Manager provides self-managing end point remediation, through patches or configuration corrections.

You can find more on this solution at:

<http://www.ibm.com/security/cisco/support.shtml>

## 7.2.2 The IBM Tivoli Security Management environment

The IBM Tivoli Security Event Management solution protects e-business infrastructures by improving response time to internal and external security threats. By centrally and automatically managing incidents and vulnerabilities from a single console, the solution provides an overall, simplified view of the security architecture. It helps administrators focus on real security threats, pinpoint hot spots, and determine the severity of attacks. Predefined tasks can help resolve threats fast, and decision support tools can quickly upgrade security policies.

The IBM Tivoli security event management solution actively monitors IT resources across an organization, filters and correlates events, and automates responses to security incidents.

IBM Tivoli Risk Manager delivers on the autonomic computing tenets of self-configuration and self-protection by assessing potential security threats and automating responses such as server reconfiguration, security patch deployment, and account revocation. Tivoli Risk Manager also produces periodic heartbeats used by upstream servers to verify that specific systems are still operational. This helps business minimize risk and business exposure, and in some instances, restore the business to its original secure state without manual intervention.

IBM Tivoli Intrusion Manager is designed to provide a “self-protection” autonomic capability in the On Demand Operating Environment through the use of a centralized management console, advanced event correlation, and reporting/analysis. By integrating a variety of sources and combining a Web Intrusion Detection System, Network Intrusion Detection System, and DB2 Universal Database, IBM Tivoli Intrusion Manager secures a customer's environment and brings together data collection, analysis, and problem resolution into a single monitoring system.

The Tivoli Security Management blueprint, in Figure 7-1, describes the different layers encompassed by the Tivoli products.

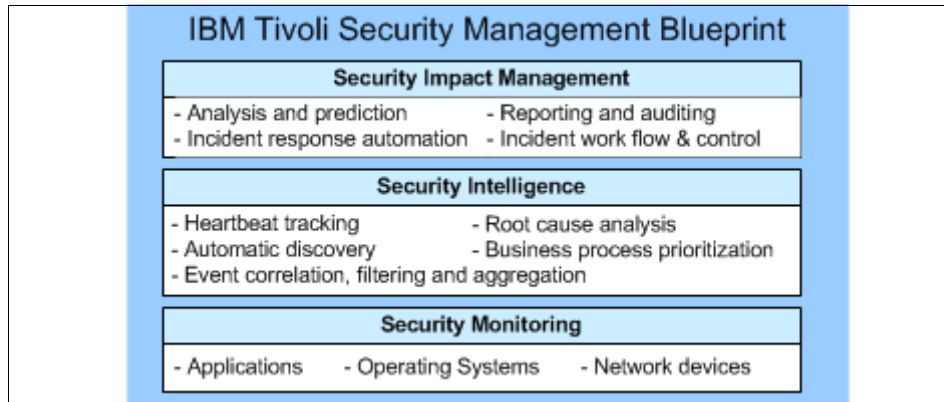


Figure 7-1 IBM Tivoli Security Management blueprint

- ▶ At the low level, we start by monitoring every key element in the security monitoring layer, which includes network devices, routers, operating systems, and applications, to make sure that everything is valid and in good shape. IBM Tivoli Intrusion Manager provides this capability by supporting a variety of sources — Checkpoint VPN-1/FireWall-1, ISS, Norton AntiVirus, Cisco Secure PIX Firewall, and Cisco Routers — with a signature-oriented intrusion detection system to bring data collection, analysis, and problem resolution together in a single monitoring system. It consolidates events, then utilizes rules and event action plans to process the events and render meaningful “situation events.”
- ▶ It also uses event action plans to store event data in a database for Crystal report generation. IBM Tivoli Intrusion Manager provides an adapter for Host IDS that can be deployed on secured operating systems to strengthen operating system security, without compromising the system's features and functionality. It provides the Network IDS sensor, which maps alarms into IBM Tivoli Intrusion Manager events to detect network-based attacks. It also includes the Web IDS sensor that detects Web server attacks. When Web IDS discovers attacks, it generates alarms, maps these alarms into intrusion-detection events, and forwards these events to the IBM Tivoli Intrusion Manager Server.
- ▶ Moving up to the security intelligence layer, IBM Tivoli Risk Manager software includes IBM Tivoli Enterprise Console and IBM Tivoli NetView to deliver a solid event correlation solution. Integration between Tivoli Enterprise Console and the Tivoli NetView gives security administrators the capability to drill down to the network topology to see where the affected resources are located and delivers true root cause problem determination. The capability to drill up from the network console to determine other resources that might be affected as a result of a network outage is also available.



The Tivoli Data Warehouse is also included in IBM Tivoli Risk Manager to store security events for long-term persistence. Tivoli Data Warehouse can store events even from other sources, such as configuration and monitoring. New data warehouse reporting capabilities can be used to leverage the information.

- In the top security impact management layer, when combined with Tivoli Business Systems Manager and Tivoli Service Level Advisor, a powerful solution capable of managing system and network resources as well as business processes can be deployed.

The intrusion detection adapters from multiple sources look for anomalies such as viruses, hacker break-ins, denial of service attacks, unusual audit log items, unusual activities, and so on. These abnormal events are sent from the adapters to the Risk Manager correlation server.

The Risk Manager correlation server analyzes information from network, systems, and devices with the sole purpose of detecting outages, threats, and unauthorized use of resources and information which are deemed critical to the success of business and, more importantly, to secure and safeguard the confidentiality of customers and their respective assets.

Consolidated events are then sent to the Tivoli Enterprise Console server for display and integration with NetView. NetView analyzes the information using a pattern matching technique, detects and predicts what is occurring based on the information collected and the pre-defined business requirements, alerts the security administrators, and helps automate corrective actions. By doing this, administrators can focus on real security threats, pinpoint hot spots, and determine the severity of attacks.

For example, there may be more than 600 messages from Tivoli Enterprise Console, but after correlation, only one event per real attack will be displayed because those 600+ messages are being consolidated and analyzed by Tivoli Risk Manager correlation server to generate the “real” security alert.

The Tivoli Enterprise Console server also has intelligence to allow administrators to define “exception” rules that certain pre-defined unusual behavior can be safely ignored. For example, administrators can define that the “server unavailable” alert generated from Application A can be safely ignored during the regular maintenance window. In this way, false alerts will be eliminated.

In a nutshell, IBM Tivoli Intrusion Manager provides adapters in various resources to detect and monitor unusual behavior and threats; Tivoli Risk Manager is responsible for the correlation and consolidation of information; NetView acts as the analyzer; Tivoli Enterprise Console is the centralized management console; and historical data is passed to Tivoli Data Warehouse.

With the Tivoli Security Event Management solution, businesses can enjoy the benefits of:

- ▶ Increasing infrastructure availability by automating problem identification and root cause analysis, leveraging best practices for problem resolution, and deploying new devices and implementing new policies quickly and easily.
- ▶ Enhancing existing security infrastructure by leveraging existing investment in security to discover problems early and take steps to resolve them, discovering security threats and exposures by analyzing historical trends, and monitoring real-time security events and identifying security threats against normal activities.

IBM Global Services Managed Security Services is an integral part of IBM Security and Privacy Services end-to-end solutions. Features include vulnerability scanning services, intrusion detection services, firewall management services, and incident management services. These features help a business evaluate current security, detect misuse and violations, respond to incidents, and implement changes to improve overall defenses.

### 7.2.3 z/OS environment solutions

z/OS introduced in V1R2 an enhanced real-time host-based Intrusion Detection Services function, using policy control to identify, alert, and document suspicious events and assist in later analysis. Host-based intrusion detection relies on additional capabilities in the host TCP/IP stack to analyze the received IP packets against known intrusion characteristic patterns. Messages about possible security violations can be sent to a log file and also sent directly to the console.

The z/OS IDS is configured with an IDS policy that defines the intrusion events to monitor along with the actions to take. The IDS policy is a set of definitions, entered as attributes in an LDAP directory, that describe to the TCP/IP stacks what events to monitor, which criteria to apply for issuing alerts, and when to trigger logging and preventive actions. The supported intrusion event types are scan detection (for example, TCP port scan, UDP port scan, ICMP scans), attack detection and prevention (for example, malformed packet events, outbound raw restrictions, inbound fragment restrictions, IP option restrictions, ICMP restrictions), and traffic regulation (such as TCP total connection and source percentage management by port, connection limiting).

In z/OS V1R5, available since March 2004, IDS is extended to provide the z/OS customer with the ability to detect misuse of resources, including the following features:

- ▶ New Intrusion Detection Support, including Physical Interface Flood Detection, is designed to detect packet floods consisting of ICMP, UDP, TCP, and unsupported protocols.


- ▶ z/OS Communications Server can work together with Tivoli Risk Manager to provide “enterprise” IDS management by collecting and correlating events from IDS sensors, now including z/OS Communications Server intrusion events, throughout the network and with an enterprise view of IDS events.

## 7.3 Key products to start with

The following IBM products should be evaluated for their use in developing a solution using this approach:

- ▶ IBM Tivoli Intrusion Manager
- ▶ IBM Tivoli Risk Manager
- ▶ IBM Tivoli Security Event Management
- ▶ IBM Tivoli Enterprise Console
- ▶ IBM Tivoli Security Compliance Manager
- ▶ IBM Tivoli Provisioning Manager
- ▶ z/OS Intrusion Detection Services functions





## **How to monitor systems to allow establishment of business SLAs and automate detection and remediation of violations**

Information technology departments are typically accountable for the service they provide for the lines of business. IT managers and line of business managers generally use service level agreements to ensure business objectives are met.

IT departments often manage service levels from multiple lines of business in a complex environment across countries and regions, over thousands of machines running on heterogeneous platforms. Without effective management tools, IT departments must use manual efforts to provide and report their services. Reporting on SLAs is always a time consuming and labor-intensive process. It requires manual comparison between data on multiple IT infrastructures, and the customers' SLAs. In addition, the IT personnel compiling reports have to deal with many different tools and consoles to find the state of resources and the business processes they affect.

Line of business managers want to know the status of the service level that may affect their business. They are not interested in the individual components of the IT infrastructure. However, that is exactly what traditional systems management tools focus on: servers, routers, applications, disk space, and so on. They do not show which problem has the highest business priority or what areas of the business are affected if a particular component is down or performing poorly.

*Service Level Management (SLM)* is an iterative process that involves identifying a business process, providing services upon which the business process depends, and defining the parameters and agents to measure the process. It is the discipline of measuring, reporting, and managing the quality of a service supplied to the business; it involves negotiating the service level agreement, monitoring the business process, fine-tuning business practices and infrastructure, and delivering increasingly better service.

## 8.1 Vision

Support for defining, monitoring, enforcing, and negotiating service level agreements will continue to evolve into a policy-driven model. Service level agreement specifications will continue to evolve and support metrics and enhanced capabilities to facilitate interactions between service providers and consumers. There will also be support for service level agreement negotiations both internal to middleware systems (for example, WebSphere negotiating with DB2) and externally.

Accessing services and information regarding them via the *Enterprise Service Bus* will enable dynamic publishing and finding of service providers meeting SLA requirements. Standards will allow for common interfaces to access this information.

Self-managing autonomic systems will integrate service level agreement awareness across all aspects of IT and business components, including process lifecycles as well as management disciplines such as configuration management, security management, workload (performance) management, availability management, capacity management, and problem determination.

The ability to map service level agreements automatically to various levels of abstraction — ranging from high-level business and IT solution infrastructures to component and system-level service level objectives — will be key to automating the process.

## 8.2 How to get started today

All operational aspects of the On Demand Operating Environment and all of its provisioned components need to be managed, as automatically as possible, based on a variety of service level objectives (SLO), service level agreements (SLAs), policies and rules.

Business Service Management provides the capabilities to visualize the IT environment in business service terms and to generate, track, and report on service level objectives. This component tightly correlates with the Policy-Based Orchestration component and the key automation capabilities: availability, security, optimization, and provisioning. With these core capabilities in place, IBM has been able to focus on building applications that take advantage of these pillars to provide true business systems management solutions. The following sections describe the functions and components which address this aspect.

### 8.2.1 The IBM Tivoli Business Service Management solution

The IBM Tivoli Business Service Management solution provides the power to align daily operations management with business priorities, set and meet service level commitments, and implement predictive management capabilities across e-business infrastructure. It works in concert with management solutions that may have already been deployed in the environment — whether from Tivoli software or other vendors — to help a business improve its ability to meet service level commitments and proactively manage e-business infrastructure.

With this end-to-end set of solutions built on a common foundation, enterprises can manage the ever-increasing complexity of their IT infrastructures with reduced staff and increased efficiency.

Mission-critical business systems typically span host and distributed environments. They comprise many interconnected applications and components, both commercial and custom, and rely on diverse middleware, databases, and supporting platforms. Tivoli Business Service Management can bridge these disparate systems to deliver a unified management solution.

The high-level integration and relationship of the Tivoli Business Service Management solution is shown in Figure 8-1.

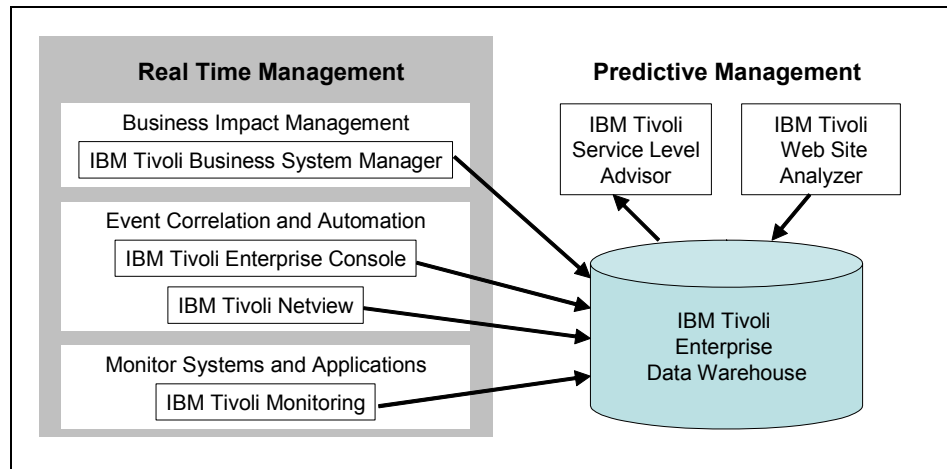


Figure 8-1 Business Service Management products integration

- ▶ On the left side of the figure, Real Time Management is split into three independent layers that offer three distinct types of value, yet provide superior management capabilities when used together.
- ▶ In the bottom layer, IBM Tivoli Monitoring monitors the hardware and software and provides automated corrective actions whenever possible to avoid critical performance problems by proactively recognizing, isolating, and repairing problems early, before these problems impact customers and other end users. It provides a comprehensive display of a transaction's path through the component systems, including response time contributions for each step. This layer focuses on the performance and availability of an individual component. From this layer you can further drill down toward the application layer where the TMTP and EWLM components of the IBM Virtualization Engine Suite can provide additional monitoring information.
- ▶ The middle layer is event correlation and automation. When problems occur that cannot be resolved at the monitoring level, event notifications are generated and sent to the correlation engine IBM Tivoli Enterprise Console and IBM Tivoli NetView. The correlation engine at this point can analyze problem notifications (events) coming from multiple components and either automate corrective actions or provide the necessary information to operators who can initiate corrective actions. The second layer helps to understand how a single failure may cause problems in related components. For instance, a router that is down could cause database clients to generate errors if they cannot access the database server.



- ▶ The top layer in real-time management is the Business Impact Management tier. IBM Tivoli Business System Manager provides insight into how a component failure may be affecting the business as a whole. For example, it helps the administrator understand exactly what line of business applications will be affected and how to reduce the impact of that failure on the business.
- ▶ On the right-hand side is Predictive Management. IBM Tivoli Service Level Advisor automatically analyzes service level agreements and evaluates compliance while using predictive analysis to help avoid service level violations. It provides graphical, business-level reports via the Web to demonstrate the business value of IT. By utilizing the automatically generated “at-a-glance” dashboard, IT staff can effectively communicate service delivery information to internal line of business executives, as well as clients. When a trend toward violation occurs, IBM Tivoli Service Level Advisor can send alerts to the IBM Tivoli Enterprise Console, to the IBM Tivoli Business Systems Manager console, or generate alerts via e-mail or SNMP.

By taking a proactive, predictive approach to service level management, IBM Tivoli Service Level Advisor allows for the analysis of trends in order to predict a service level objective violation before it occurs. Taking corrective actions to prevent SLA violations keeps improves customer satisfaction and overall business performance.

IBM Tivoli Web Site Analyzer uses Web server data to identify performance trends and visitor traffic patterns in a way that can help to evaluate the impact of Web-based programs.

All the historical system management data gathered from IBM Tivoli Monitoring, IBM Tivoli NetView, IBM Tivoli Enterprise Console, IBM Tivoli Business System Manager, IBM Tivoli Web Site Analyzer, and other sources is stored in IBM Tivoli Data Warehouse for long-term persistence.

Chapter 14, “How to monitor end-to-end applications, their topology, and their resources” on page 159 and Chapter 19, “How to monitor using EWLM” on page 295 provide one aspect of monitoring that may complement this approach, using Enterprise Workload Management.

## 8.3 Key products to start with

The following IBM products should be evaluated for their use in developing a solution using this approach:

- ▶ IBM Virtualization Engine Suite (EWLM Component)
- ▶ IBM Tivoli Access Manager for Business Integration (prior to March 2003, this product was known as IBM Tivoli Business Systems Manager)
- ▶ IBM Tivoli Business Systems Manager for z/OS
- ▶ IBM Tivoli Monitoring
- ▶ IBM Tivoli Service Level Advisor
- ▶ IBM Tivoli Web Site Analyzer



## **How to reduce the time and cost to re-purpose IT resources to meet business requirements**

In today's business environment, companies are constantly under pressure to control or cut costs. On Demand Business is a strategy that can be applied to realize benefits and value across the entire enterprise. It is not something that needs to be focused only on production applications and resources. Companies can achieve savings and value by exploiting on demand to re-purpose IT resources for other areas of the business.

One important characteristic of an On Demand Operating Environment is its ability to adapt and configure itself to changes in the infrastructure, with minimal human intervention, based on goals and policies specified by the offering or provider administrator. Some triggers that drive configuration changes in an On Demand Operating Environment can include an administrator's requests to allocate/de-allocate or reconfigure resources; requests from other autonomic managers (for example, a workload manager, availability manager, and so on) to allocate/de-allocate resources or reconfigure resources supporting an already running service; and dynamically detected configuration changes that have implications to other resources or other resource managers.

In all of these cases, the actual changes that need to be made to the environment could range from simple changes (which involve the execution of a single command), to more complex changes that need to be orchestrated across multiple resources and resource managers. Typically, such configuration changes tend to be labor-intensive and error-prone tasks that require documentation of and adherence to complex change-management processes.

This approach is about providing adaptable capabilities to automate the management of IT resources, including rapid, error-free, and less labor-intensive re-purposing of IT infrastructure elements to meet business and application requirements.

## 9.1 Vision

To provide these capabilities, IBM provides an architecture approach as described in the following paragraphs.

### 9.1.1 Architecture approach: the MAPE Loop

The IBM Autonomic Computing technology initiatives have put back to the foreground a classical retro-feedback loop to apply it into IT environments. Part of automation, this loop describes all the needed steps for an infrastructure management product to apply changes into the environments. As described in Figure 9-1, this loop is commonly called the *MAPE Loop*, due to the names of the four steps used: Monitor, Analyze, Plan, and Execute:

1. Use sensors to collect information through the *Monitor* entry. Example:

%CPU of box A= 97%  
Response time of box A= 750ms

2. *Analyze* this data to generate a result (a kind of satisfaction criteria) to be compared with a predefined objective. Example:

CPU Usage is too high (I would expect it below 90%)  
Response time is bad (I would expect it below 300ms)

3. If this objective is not met, one solution is *Planned* to correct it. Example:

I will try to add a new CPU, but where? In the same box, in a new box?  
...  
Box A has no more CPU, I only have new boxes available.

4. Apply this correction with the *Execute* step, by using effectors in the environment. Example:

I will prepare box B to do the same needed workload: (OS, Products, Applications, Network).

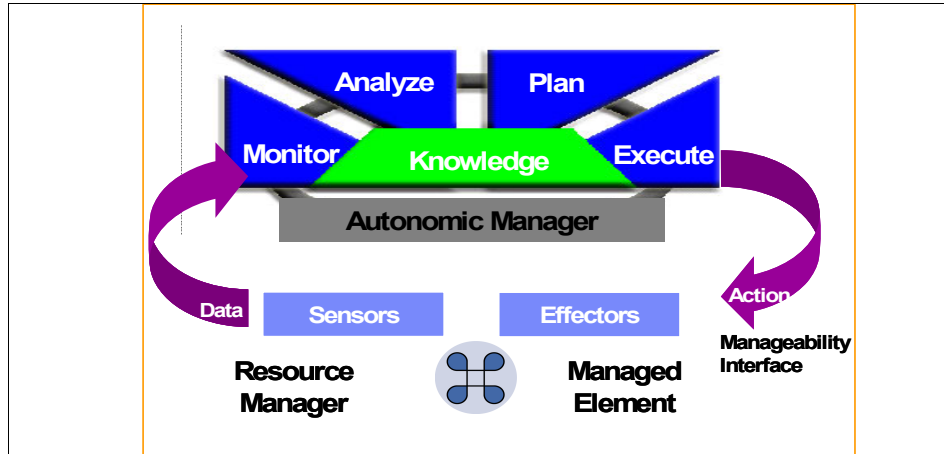


Figure 9-1 The Autonomic Computing MAPE Loop

Orchestration and provisioning solutions will be expanded to provide provisioning support to encompass more applications, middleware, operating systems, and networks. This will be enabled through standards such as the Open Grid Services Architecture (OGSA). This combination of an IT-level SOA and OGSA-based interfaces allows for dynamic provisioning of virtualized resources in an heterogeneous environment.

The ability to reduce the time and cost to re-purpose IT resources to meet business requirements and service levels will evolve to an autonomic assignment of resources based on business policies and self-optimization. The implication here is the continued evolution of OGSA and autonomic functions facilitating the integration and management of all components making up the business solution.

### 9.1.2 Optimizing existing resources

Grid solutions optimize server utilization and leverage under-utilized resources: they participate in workload management, can use excess capacity during peak off periods or standby resources when they are not used for disaster recovery or development and test. In other words, grid solutions optimize use of IT investments by creating an environment that enables sharing of under-utilized computational resource. These methods have been used for many years in scientific/technical computing; combining the possibilities of Internet technologies with the facilities provided by newly adopted open standards will spread the use of these grid technologies in commercial environments.

### 9.1.3 Storage approach

Storage virtualization products, with the non-disruptive data migration features, strongly reduce the time to re-organize the storage infrastructure by avoiding the disruptive backup/restore procedures that are generally required to move data between different physical storage subsystems.

## 9.2 How to get started today

The following sections describe the different components to set up the solution.

### 9.2.1 The MAPE Loop implementation

Starting from the MAPE Loop, we can easily position all the useful products. TotalStorage products that solve this infrastructure management issues are also described.

#### **The “Monitoring” components**

*IBM Director Multi-Platform* (Director MP) provides resources monitoring, based on Common Information Model (CIM) entries. These objects are triggered by thresholds to give a status (OK, warning, or critical), and can be caught by the Event Manager. The Director bridge is a VE component to link Director information, status, and tasks to the VE Console.

*IBM Tivoli Monitoring* provides a wide monitoring scope and extends its monitoring capabilities to non-IBM OS and machine. It can also be linked to the VE Console.

*IBM VE Console* is a portal-based application, used to create a single point of administration. The console maps several types of tasks, such as resource management and monitoring. Hence the console can offer to a specific user a mix of several various tasks to limit the console numbers and the action scope of this user. All these tools and information collected at the VE Console level are coming from management sources. These kinds of plug-ins are reachable through Web Service connections and make available different types of data, coming from different products. There are four management sources today:

- ▶ IBM Director MP, through a VE bridge
- ▶ IBM Cluster Systems Management, through a VE bridge
- ▶ IBM Tivoli Monitoring, with a direct connection
- ▶ IBM Management Central, through a bridge

*IBM Enterprise Workload Manager* has a monitoring part to collect data from ARMed middleware and to classify defined transactions. This end-to-end monitoring gives a correlated view between all the tiers. The EWLM agents collect this data; they are gathered at the EWLM management server and are accessible via the Web based graphic interface.

### **The “Analyzing” components**

*IBM Enterprise Workload Manager* has also an analysis component to validate the achievement of different objectives which has been deployed in the active policy. The data, collected previously, will be used to determine the *performance index* of each classified transactions. This measure is the starting point to decide what to do if the objectives are not met. EWLM will have its own *plan* component or will use another product, but it is not defined at the time of writing this book.

### **The “Planning” components**

*IBM Tivoli Intelligent ThinkDynamic Orchestrator* (TIO) extends the benefits of the IBM Tivoli Provisioning Manager (TPM, see below). It orchestrates the activities necessary to automatically maintain server availability and meet required service levels. It provides the *why*, *where*, and *when* of a complete solution. By monitoring the applications under its control, TIO can sense degrading performance and determine why actions need to be taken. Because solutions are monitored closely, TIO can determine where (for which application) a resource is needed and instruct TPM to deploy a server, install the necessary software, and configure the network. This enables an application to maintain acceptable service levels.

The idea behind this approach is to provide an IT infrastructure environment in which the needs of a business process can be met by having the required resources provisioned out of a common pool of shared resources. This common pool of resources, which could be servers and storage, for example, would be managed and controlled by TIO. By working together with the TPM, it would support each business application's need to have servers provisioned and built for them as they require resources. This approach allows the same resources to be re-purposed if the system requirements are reduced. For example, when not required, the servers would be de-provisioned and placed back into the common resource pool, where they would be re-purposed and made available to the next requestor of these types of resources.

The process of provisioning these resources is automated to minimize the amount of effort and time needed by skilled IT resources to rebuild and reconfigure the environment for each line of business. The work gets done faster, with less skilled resources, and it gets done consistently. This minimizes the amount of errors that might normally occur during a manual build and improves the quality of the deliverable by the IT organization.

Accomplishing this requires planning, and guidelines need to be established with the line of business teams to ensure they understand how the environment will provide them with the service levels they need as well as help in cutting their direct costs and the costs to the company. The right amount of resources that must be kept in the pool also needs to be determined based upon the projected need for those resources.

Also, the frequency and speed at which these environments need to get built depends on the needs of each business process. Knowing the speed at which they need to be built will help the IT organization determine whether certain resources are pre-loaded with a certain operating system or certain piece of application middleware, for example. This would apply if all applications were standardized on a certain operating system such as Linux, or on a specific middleware application server or DB such as the IBM WebSphere Application Server and IBM DB2.

### **The “Executing” components**

*IBM Tivoli Provisioning Manager (TPM)* provisions and configures servers, operating systems, middleware, applications, and network devices acting as routers, switches, firewalls and load balancers. It provide the *how* and *what* of a complete solution. TPM, through workflows, automates the manual provisioning and deployment process. Using pre-built “industry best practice” workflows to control and configure popular vendor products, or user-customized workflows to implement environment or proprietary application-specific “best practices,” processes can be automated and executed in a consistent error-free manner.

TPM needs to know all the existing infrastructure. The description of all components is essential to work on them. Networks, switches, servers, OS, software, and applications need to be described. All of this data is collected and organized in the Data Center Model (DCM). The DCM is very critical to create procedures which will call or modify these components. In TPM, procedures are called workflows.

As we discussed previously, TPM provides pre-built workflows. They allow you to have a starting point to use them in more complex workflows, or to customize them. These are the predefined steps that need to be performed in order for the resources to be built and deployed. The workflows provide the repeatable best practices and company-specific policies for the automated portion of the solution to build and deploy the servers. These workflows are used to build automation packages whereby the workflows can be reused. This would occur when, for example, a Web application server is installed. Certain core steps would be performed to complete the installation of the software. This process could be reused for installing the same Web application server code on other machines for another developer team.



The last important feature in TPM is the component group on which workflows will be applied. They are named drivers. These drivers represent all the potential exit-point from TPM to modify remotely a configuration. Here are a few examples of these drivers:

- ▶ SNMP drivers, to apply changes in a network configuration
- ▶ SSH drivers, to execute scripts on a machine
- ▶ NIM drivers, RDM drivers, to install images in a pSeries or xSeries box
- ▶ VMware and zVM drivers to create or modify Linux guests

IBM provides automation packages that deploy software images onto servers as well as other provisioning functions. The workflows interact with existing IBM @server products to automate software deployment under the control of TPM or TIO.

TPM knows what resources are available and what resources are needed to deploy a specific application server. Once it knows what software must be loaded onto those configurations, and how to orchestrate the installation and customization of that software, TIO does the coordination between its various subcomponents to provision the environment and track the resources assigned to that configuration. The provisioning and de-provisioning processes can also be triggered manually by an administrator when requested by a line of business.

*IBM Enterprise Workload Manager* also owns one “execute” component. EWLM can act on CISCO and Nortel load balancers to dynamically change the weights of request routes. This effectors can be triggered after the analysis of the classified transactions.

## **MAPE Loop summary**

Implementation and exploitation of all these products within this approach can provide the following benefits:

- ▶ Reduce costs and improve the return on investment associated with equipment and software needed by the enterprise by re-purposing those resources.
- ▶ Reduce the number of skilled personnel and time required to build and configure the environments by automating provisioning and configuration.
- ▶ Improve the quality of the built configurations by using proven, repeatable automated workflows and processes.

Besides the cost savings that can be realized by sharing the servers among the various lines of business, additional savings may be possible by re-using software licenses. For example, if the company is licensed to run a certain product on a 4-way Intel processor, that same license can normally be used to satisfy any user of the system as long as it is only deployed on that type of server, and used on only one server at any given time.

## 9.2.2 The role of the grid

All IBM @server platforms provide a path to grid: the Globus toolkit is available on all platforms, and grid ISVs support multiple types of environment. OGSA will be packaged on all platforms. Each platform provides a unique grid value:

- ▶ Linux is supported across all platforms; in particular, Linux can run on the same box as z/OS, with the Integrated Facility for Linux (IFL), without being subject to software licensing fees on traditional partitions.
- ▶ xSeries, pSeries, and iSeries for, respectively, Intel workloads and blade packaging, RISC workloads and cluster facilities, and OS/400 and partitioning.
- ▶ zSeries for z/VM virtualization supports the creation of hundreds of virtual Linux servers on a single box, can dynamically deploy and configure new virtual grid servers and add hardware capacity in just minutes, and can provide speedy access between application partitions on the same mainframe using HiperSockets. The coexistence with z/OS systems can be well managed if the weights between partitions are set up correctly. Linux grid workloads can use all the capacity of the machine not currently used, without impacting the production systems.

## 9.2.3 Storage implementations

*IBM TotalStorage Productivity Center for Fabric* (formerly IBM Tivoli Storage Area Network Manager) has two built-in features to help re-purpose IT resources to meet business requirements. First, a zone control feature that allows zones to be created and maintained on the different switches of the fabric from a central point. Then it is also able to provide SAN fabric performance and capacity management reporting and monitoring with a graphical display.

Both *IBM TotalStorage Productivity Center for Fabric* (formerly IBM Tivoli Storage Area Network Manager) and *IBM TotalStorage Productivity Center for Data* (formerly IBM Tivoli Storage Resource Manager) are designed to allow a system administrators to leverage their knowledge and skill across the enterprise to more effectively manage the storage assets of the company.

Their architecture enables the system administrators to see in a graphical way all of the storage assets, including direct attached storage (DAS), network attached storage (NAS), and storage that resides on the storage area network (SAN). This comprehensive view of the entire storage map allows the administrators to manage much larger environments, but also get the information about utilization that is typically required in large environments.

IBM TotalStorage Productivity Center for Fabric, which provides a centralized zone control feature, allows zones to be created and maintained on the different switches of the fabric. It includes also SAN fabric performance and capacity management reporting or monitoring feature. These two build-in feature can also help to reduce time and cost when storage resources need to be reorganized.

Data collection by the IBM TotalStorage Productivity Center for Data occurs automatically on a customer-defined schedule with the resulting data residing in a repository of storage metadata. The customer's storage metadata can then be extracted with a graphical user interface that is invoked as an application or as an applet via browser, or the data can be automatically published in a number of formats including HTML for storage portal presentations.

The information collected by the IBM TotalStorage Productivity Center for Data allows the customer to make intelligent decisions to reduce the waste space that is usually found in open systems environments; it allows to take decisions and actually steps to increase effective utilization of their storage environment. The scope is not just limited to files and their attributes, but it addresses also the waste space that is found within relational database managers such as Oracle, Sybase, SQL/Server, and DB2.

## **SVC Implementation examples**

The virtualization layer implemented by the SVC in the storage infrastructure provides a physical storage abstraction view; using advanced migration features, the storage infrastructure can then be reorganized non-disruptively for the applications<sup>1</sup>. For example:

- ▶ The storage pools, from which the virtual disks are created, can be transparently reorganized at any time: for example, if a storage pool needs more disk space, some other free physical disks (LUNs) can be dynamically added to that pool; thus some storage space can be allocated for new virtual disks creation.
- ▶ If LUNs, used in a storage pool, need to be moved to an other storage pool, the SVC allows this LUN to be migrated non-disruptively; data blocks residing on that LUN will be automatically redistributed on the other available LUNs of the former storage pool. This is an SVC controller feature called “Volume drain”.
- ▶ An other very useful and migration feature provided by the SVC is the “Data Migration” which allows dynamic virtual disk migration between storage pools without impacting applications.

---

<sup>1</sup> To be accurate, groups of LUNs with SVCs are called “Managed Disk Groups” and these groups are called “Storage Pools” with SFSs; LUNs are called “Managed Disks” by SVCs and “Volumes” by SFSs. To simplify as follows, we used the words “Storage Pools” for both SVCs and SFSs.

With this major feature, a virtual disk defined in a storage pool using ESS LUNs can be moved to an other storage pool using other LUNs, from other storage subsystems (IBM or non-IBM). For example, if some data becomes less critical, it will be possible to migrate the corresponding virtual disks on a less expensive storage pools (storage subsystems). Without that feature, data migration between different storage subsystems requires to go through standard backup/restore processes, implying an application down time.

### **SFS implementation examples**

The other IBM storage virtualization product, SFS, provides also interesting features such as file movement, which helps in storage infrastructure re-organization. With SFS, file created in the Global File system are physically located in a specific storage pool and uses a policy based file placement mechanism. A policy is a set of rules that determines where the specifics files have to be placed in the storage pools, based on the files attributes (type of file, date of creation, file size, user ID, etc.).

Once created and used by SAN File system clients, it is still possible to move dynamically the files to other storage pools if necessary.

- ▶ For example we can use a file movement command that moves one or more files from their current storage pool location to a different specified storage pool.
- ▶ It is possible to automate these file movements by using a policy-based file management mechanism. This feature provides life cycle management capabilities to SAN File System. A file management policy is also a set of rules, which specify conditions for either automatically moving files from one storage pool to another, or for deleting them entirely from the storage pool.

Both **SVC** LUNs or Virtual disks migration and **SFS** file movement feature could help each time storage infrastructure reorganization is necessary in response to business changes requests with a minimum time, cost and impact on application (all migration or file movements require each only one SVC or SFS command to be executed).

## 9.3 Key products to start with

The following IBM products should be evaluated for their use in developing a solution using this approach:

- ▶ IBM Tivoli Intelligent ThinkDynamic Orchestrator
- ▶ IBM Tivoli Provisioning Manager
- ▶ IBM TotalStorage Productivity Center for Fabric
- ▶ IBM TotalStorage Productivity Center for Data
- ▶ SVC
- ▶ SFS
- ▶ IBM Grid Toolbox
- ▶ IBM zSeries processors with LPAR
- ▶ IBM z/VM and Linux for zSeries
- ▶ IBM Virtualization Engine components:
  - Enterprise Workload Manager
  - IBM Director Multiplatform
  - IBM Virtualization Engine Console





## How to map IT resources used by various business processes of an end-to-end solution

Organizations need a complete solution that helps IT specialists efficiently track assets from a financial, contractual, and usage point of view. An integrated view of its software assets can help an organization effectively plan for maintenance and upgrades and understand precisely which resources are needed to support their business.

Similarly, an understanding of the components that make up a solution, and their relationships, is critical to managing the solution in its entirety as well as managing the various components in the context of the solution (versus managing them as independent elements).

This understanding allows the decomposition of the performance of a solution into various tiers. An autonomic manager can then invoke appropriate configuration changes or provision additional resources where they will have the most impact towards meeting service level agreements and business objectives.

## 10.1 Vision

In an On Demand Operating Environment, where many resources are virtualized and shared between lines of business, it is difficult to associate specific resource usage with a business process or business unit. However, for charge-back purposes this is still an important capability. A relationship registry that can track allocated resources to specific business processes or applications will help meet these requirements.

In addition, the common event infrastructure will make it easier to track all events within the environment. These common events could be used to help indicate the usage of resources by specific processes and therefore map resource usage to business areas.

Of course, standards play a key role in enabling such capabilities. Some of the standards that apply and will make it easier to develop and track applications and their resource usage include BPEL and OGSA.

## 10.2 How to get started today

The understanding of the components that make up a solution and their relationship is critical to managing the solution in its entirety. In today's environment, many system resources are still allocated to a specific application or business process. Therefore, we also discuss an approach for installing and tracking software and licenses to specific systems.

### 10.2.1 Understand the usage of data

The data collected by the *Tivoli Storage Resource Manager* enables customers to understand what is really going on with the data that resides on their servers. This includes views that show when files are created, accessed, and modified, and by what group or user. This type of information enables the system administrators to map the actual storage resource to the consumers of that resource. The ability to map storage consumption to storage hardware has become increasingly important as the size of the open systems environment has increased and the cost of the underlying hardware has become more expensive with the advent of enterprise-class storage in the open systems arena.

In addition to understanding the current consumption and usage of data within the enterprise, the Tivoli Storage Resource Manager keeps track of this information over time. Not only does this historical view of storage consumption and utilization allow the customer to see usage trends over time, it also enables the system administrator to see a projected use of storage into the future. This



allows the system administrator to prepare for the need to purchase additional capacity in a planned proactive manner rather than just reacting to being out of space again.

## 10.2.2 Understand the usage of specific components

Multiple tools can be used today to understand the usage of components, these resources can be hardware or software related.

- ▶ *IBM Tivoli Configuration Manager* can help gain total control over enterprise software and hardware. Its software distribution module can provide the ability to rapidly and efficiently deploy complex mission-critical applications to multiple locations from a central point. After systems have been deployed, the inventory module can automatically scan for and collect hardware and software configuration information from computer systems across the enterprise.
- ▶ *IBM Director MultiPlatform* proposes an interesting view of the infrastructure because of three features:
  - Director MP manages resources and creates groups. These groups are static (default Director group or created by the user) or dynamic (by applying filters to check the membership of a resource).
  - Asset Manager in Director can also manage the hardware layer (firmware, device drivers), but also leasing info or whatever by allowing administrators to update fields in the configuration.
  - The Software Repository can also collect data from managed machines to group machines or make queries.

The query tool can help to create dynamic groups or identify resources with specific criteria. These criteria can be any data coming from the Common Information Model (OS type, CPU, memory, network devices,...), from the asset info (firmware type, date, leasing dates) or from the software repository. It is very easy to query resources to identify machine in a specific configuration. For example, you can find in your environment all the machines which need an upgrade or a software maintenance.

- ▶ *IBM Tivoli License Manager* is a tool for managing software licensing. With advanced inventory and reporting capabilities, it helps businesses know exactly what software licenses they have and which ones they might need. IBM Tivoli License Manager enables businesses to ensure compliance and reduce costs through advanced inventory and reporting capabilities.

### 10.2.3 Sample scenario

The following scenario, based on a fictitious business, illustrates this capability.

#### Scenario background

ABC Corporation (a fictitious business) is a worldwide jewelry manufacturer with headquarters in Shanghai and subsidiary factories in San Francisco and Milan. Flora is the Vice President of the IT department. Recently, there are three major issues on Flora's mind:

1. There is an urgent need to deploy new computer-aided design (CAD) software to all of their jewelry designers' desktops so that they can be more efficient (and thus competitive) when designing their products. Due to the geographical distance between the three regions (Asia, EMEA, and US), an automatic deployment with continuing monitoring of the environment to keep the CAD software level up-to-date is highly desirable.
2. There is a need to cut expenses. Due to budget constraints, no additional resources can be allocated to deploy the new CAD software.
3. When the help desk receives trader complaints about poor response time on trading transactions, it often takes several days with several highly skilled personnel involved to identify the root cause. The troubleshooting process is labor-intensive, slow, complicated, and error-prone. This reduces traders' confidence in ABC Corporation and is affecting the company's reputation.

With an infrastructure management solution, Flora can handle the problems without interrupting ABC's business operations.

#### Phase 1: Reducing the number of licenses

First of all, Flora needs to identify the overspending areas. She does this by using IBM Tivoli License Manager to examine the actual use of ABC's most expensive application, Oracle Database Enterprise Edition. ABC has developed an in-house Web-based application to manage its trading business, and the information from this application is stored in the Oracle Relational Database Management System. ABC estimated there would be 20 developers working on the staging, testing, and production environments, and 1,000 traders would be using the Web site that resides on the production environment. So, Flora bought 1,020 Name Users Plus licenses from Oracle last year.

Based on the IBM Tivoli License Manager software examination, the actual number of developers working on the Oracle software is 12 and 600 traders are using the production Web site for the trading business. Thus 408 Oracle Name Users Plus licenses are never used.

With the figures provided by IBM Tivoli License Manager, Flora now has confidence that the 1,020 Oracle Name User Plus licenses can be cut down to 612 without any significant impact to the business operation.

A user pool is created based on the 612 user names gathered from IBM Tivoli License Manager during the software examination. This pool ensures the 612 users continue to have access to the software.

IBM Tivoli License Manager continues to enforce the license contracts by ensuring only the users that are identified in the user pool will be able to access the Oracle software. This capability helps Flora to maintain her license contracts with the software vendor and to avoid any penalty.

The 40% savings from the Oracle Name User Plus license can now be re-allocated to the new CAD software deployment with IBM Tivoli Configuration Manager for software distribution automation.

The solution design does not require specific constraints within any of the Tivoli resources. As much as possible, the solution configuration parameters are stored in a relational database. Those parameters are captured either via dialogs developed using the Application Extension Facility (AEF) or by using batch processes that capture the endpoints and software package definitions.

The process structure uses two layers: an operational layer and an automation layer. It also makes use of the Inventory database to maintain the structure needed to facilitate true automation, as well as some of the associated Tivoli administrative information. The Inventory database stores not only the software inventory information, but the software package to endpoint relationship. The use of the Tivoli Desktop interface has been retained as the management interface by the use of customized dialogs (AEF). This is to ensure as much compatibility with the out-of-the-box Tivoli products as possible. The depot functionality of IBM Tivoli Configuration Manager is used in conjunction with the Activity Planner to prepare plans that will load the depots with software packages. This allows software package blocks (SPB) to be loaded on the depot server and then installed to the endpoint at a later time.

## **Phase 2: Preventing problems with self-healing capability**

The self-healing capability in IBM Tivoli Configuration Manager is designed to be used in a per request mode more than in a massive distribution scenario. For example, we may suspect that a set of workstations might have problems with a set of software packages that were distributed using the IBM Tivoli Configuration Manager. We can assign a self-healing role to the suspect workstations, and an automatic Verify and Repair operation can be launched by the Tivoli Management Region (TMR) Server.

The *Verify* operation checks that the changes that occurred during the last operation are intact, that is, the files or any other package setting defined using the Software Package Editor actually remain intact. If this operation fails, the package is set to Error. Verify compares the missing settings versus the software package stored in the Source Host.

The *Repair* operation installs only those source system objects that are missing on the target system to make the target system objects consistent with the source system objects defined in the software package when using the Software Package Editor.

They also learned that a critical update has been released on the IBM Web site. This update covers a xSeries family firmware. Using the IBM Director query tool, they found that three machines are impacted by this fix. After downloading the fix, they used the Update Xpress tool to deploy it simultaneously, with a minimum effort.

The most critical issues of ABC Corporation are easily and quickly resolved with these software applications.

## 10.3 Key products to start with

The following IBM products should be evaluated for their use in developing a solution using this approach:

- ▶ IBM Tivoli Configuration Manager
- ▶ IBM Tivoli License Manager
- ▶ IBM Director Multiplatform
- ▶ IBM Tivoli Storage Resource Manager



## How to consolidate and simplify the IT infrastructure

As IT infrastructure is becoming more complex, solutions are required to simplify its management and reduce its cost, to improve its reliability, and to prioritize resources dynamically in order to support the business better and faster.

To accomplish these improvements, it is necessary to:

- ▶ Lower the cost of the existing infrastructure.
- ▶ Reduce the complexity of adding new elements to that infrastructure.
- ▶ Reduce the complexity of managing the infrastructure.
- ▶ Build new infrastructures that are more responsive to business needs.

Various solutions or technologies are available to simplify and optimize the IT infrastructure. These can be implemented in different areas of the infrastructure, such as the servers, the storage, the network, etc. We can identify the following technologies:

- ▶ Systems consolidation is more than just replacing many smaller servers with a few larger servers. It is about simplifying and optimizing existing end-to-end IT infrastructures, including servers, storage, databases, applications, networks, and systems management processes. The goal is to reduce both cost and complexity. Consolidation also provides a more efficient and stable foundation for growth and new solution development, which makes it the logical first step toward deploying an On Demand Operating Environment.

- ▶ Dynamic provisioning is another technology that enables consolidation of servers by allowing them to be shared across workloads, driving up utilization and allowing for a reduction of the number of servers and storage systems.
- ▶ Virtualization is yet another technology that can be implemented in many areas of the infrastructure. Virtualization provides a logical rather than a physical view of the resources such as computing power (servers), storage capacity, or other resources. This makes possible a basic system management of multiple disparate systems and allows a real-time dynamic deployment and optimization of these resources. Most of the time, this technology makes an aggregation or consolidation of the existing resources, which also simplifies the infrastructure management.

As shown in Figure 11-1, there are several methods of consolidation that can be implemented individually or in combination:

- ▶ Logical consolidation involves consolidating servers supporting different types of workloads onto fewer or larger systems.
- ▶ Centralization co-locates physical servers and storage into fewer locations to reduce the number of physical sites.
- ▶ Physical consolidation combines similar server and storage resources into larger, more powerful systems, while typically maintaining the same application platform.
- ▶ Data consolidation combines data with different formats into a similar format or platform.

Consolidation allows companies to leverage IT for added business value. As IT reduces complexity and simplifies systems management, total cost of ownership (TCO) dramatically improves. Part of the TCO improvement comes from the ability to reduce and eliminate redundancy in the infrastructure and gain new efficiencies for the business that help to maximize revenue per customer.

As the infrastructure realizes improvements in availability, recoverability, performance, and scalability, the business can respond more quickly to new challenges and provide better service to customers, employees, suppliers, and partners. Through this consolidation, specific efficiencies may be realized, such as freeing up support staff to perform other tasks and reducing space and energy requirements.

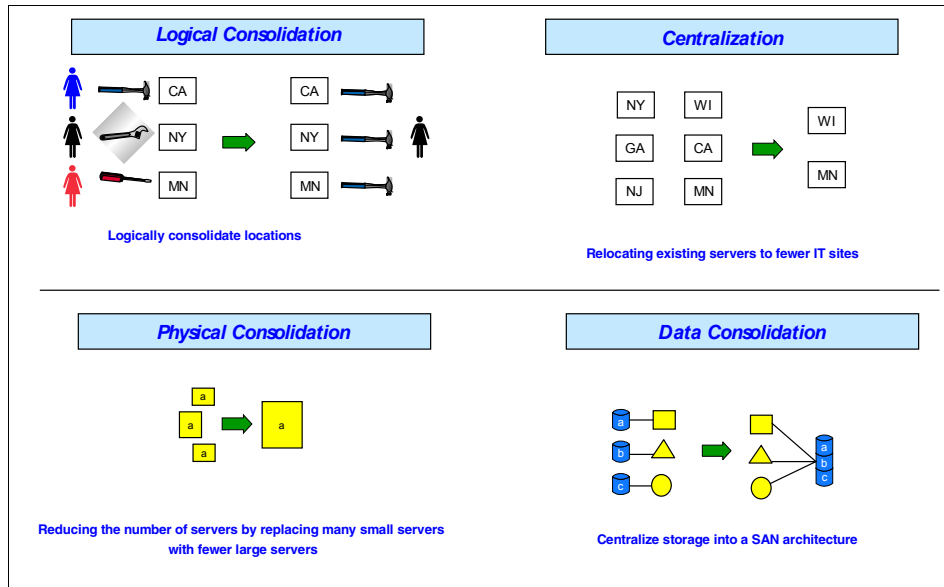


Figure 11-1 Types of consolidation

Considering a storage infrastructure, for example, we can describe physical and logical consolidation as follows:

► **Physical consolidation:**

Data from disparate storage subsystems can be combined onto large, enterprise class shared disk arrays, which may be located at some distance from the servers. The capacity of these disk arrays can be shared by multiple servers, and users may also benefit from the advanced functions typically offered with such subsystems. The array capacity may be partitioned, so that each server can use an appropriate portion of the available gigabytes. Available capacity can be dynamically allocated to any server requiring additional space and capacity not required by a server application can be re-allocated to other servers. Extra capacity may be added, non-disruptively. However, physical consolidation does not mean that all wasted space concerns are addressed.

Physical storage consolidation is mainly implemented into the storage subsystems themselves, whose main role is to provide logical disks, called Logical Unit Numbers (LUNs), to different servers that can run different operating systems and different applications. These storage subsystems may include RAID capabilities, remote mirroring, and instantaneous data replication functions, which might not be available with smaller integrated disks.

Now we can say also that these storage subsystems already implement a kind of virtualization because the LUN, which are seen as disks (or SCSI targets) by the servers, are in fact virtual disks; they have been created by the storage subsystems from the available space on its physical disks.

Consolidation of storage volumes in few bigger storage subsystems automatically simplify the IT infrastructure (and its management) by reducing the number of its elements. Each new generation of storage subsystems (IBM DS4000, DS6000 or DS8000 family) is more powerful in terms of throughput and storage volume, as well as features (copy services) that automatically improve their consolidation capacity.

The importance of these highly reliable and high performance storage hardware solutions as the guardian of mission-critical data for the business is still a cornerstone concept. However, software is emerging as a critical element of any SAN solution; management and virtualization software provide, by their concepts, advance functionality for administering distributed IT assets, maintaining high availability, and minimizing downtime.

► **Logical consolidation:**

With logical consolidation, it is possible to allow any application server to use any storage system available space. With logical consolidation, logical disk resources may allow available capacity to be allocated and re-allocated between different application servers in a non-disruptive way, allowing a better utilization of the full storage space. Logical consolidation will be generally implemented by storage virtualization solutions.

Storage virtualization solutions make a logical consolidation into the IT infrastructure when they are able to combine several storage subsystems in a single storage reservoir from which virtual elements such as Virtual Disks or Storage Pools could be created. This storage consolidation (or aggregation) can be homogenous or heterogeneous (multiple vendors).

The storage virtualization solution management tool becomes also a central management point, allowing basic systems management of multiple disparate systems, which simply the IT storage infrastructure management itself.

Storage virtualization solutions provide a logical view rather a physical view of the storage, allowing the implementation of advanced features such as policy based automatic data placement or data movement, data migration, and policy based Life Cycle Management. These capabilities can strongly contribute to the storage infrastructure management simplification.



## 11.1 Vision

From an infrastructure perspective the different consolidation scenarios may be distinguished by the operation environment requirements. The scenarios of physical or application consolidation can be refined to either full or virtual consolidation:

- ▶ With full consolidation, several servers, running either the same or different applications, are consolidated on one server and one operating system image.
- ▶ With virtual consolidation, servers are consolidated into the same number of operating system images running on one box by using a virtualization technology.

There are several reasons why a full consolidation may be difficult: security, stability, maintenance, flexibility, or support. With virtual consolidation, resource allocation is secure, transparent, and fully dynamic.

Businesses will be able to consolidate multiple operating systems across a common hypervisor. To be precise, a *hypervisor* is a firmware component upon which LPAR and DLPAR are built. With this capability, fewer physical systems can be used, thereby reducing floor space and energy requirements. In addition, it enables the dynamic incorporation of operating systems that can communicate with each other in a secure, reliable, and fast way. Sub-processor logical partitioning, already available on some platforms, will expand and be available on additional platforms. This provides fine-grained control over the resources allocated to a specific application or process.

Consolidation, tied with autonomic provisioning techniques, will provide yet another level of IT optimization across the organization (imagine being able to automatically provision and de-provision servers on the fly during peak business hours).

Workload management is a critical component of server consolidation. Through grid computing initiatives, especially with respect to job scheduling and the allocation of work to specific virtual systems, economies of scale can be gained by running applications on partitioned servers.

By moving intelligence into the storage network with storage virtualization products, storage control is moved also into the storage network, which offers the opportunity to reduce complexity by providing a single view of the storage.

# 11.2 How to get started today

Many people think of physical consolidation when they first consider server consolidation. Physical consolidation is the process of reducing the total number of servers by replacing many servers with fewer, more powerful ones, or clustered systems. This can take place within the same architecture or across architectural boundaries.

The first step in server consolidation was brought about by the design of more and more powerful systems. This is true for every brand. Take the zSeries 990, the pSeries 690, the iSeries 890 or xSeries 445, as examples. All of these servers now have the power of several servers just 2 years ago. But power is not enough if it can't adequately be shared, and this is where new capabilities make server consolidation a great cost-saving solution. Some of the technologies that answer the need for sharing are defined in Table 11-1.

Table 11-1 Definitions of consolidation technologies

Technology	Definition
Hypervisor	The hypervisor is the firmware component which allows the split of resources between different partitions. Its role is to make sure individual resources within the system are only accessed by the authorized operating systems.
LPAR	Logical partitioning is the ability to run several instances of operating systems on the same hardware systems. We can distinguish several steps in logical partitioning. Dynamic Logical Partitioning (DLPAR) allow the resources to be moved from one operating system to another without having to reboot the OS.
Virtual machine	A software counterpart to LPARs, this is also a way to create a larger number of operating system images within an already running operating system, (rather than a hardware platform for LPAR). Virtual machines are available on zSeries, with z/VM, or on xSeries, with products such as VMware.
Workload Management	More than just load balancing, Workload manager classifies the incoming workload according to the installation policy and makes sure that each of the different workloads will be able to use the amount of resources that they require in order to achieve the goals set in the policy. From a server consolidation standpoint, it enables the ability to go from standalone machines to a consolidated environment with no degradation in terms of performance.

## 11.2.1 zSeries implementation example

Of course, none of these technologies are exclusive. If we take the zSeries brand, for example, we have the four layers coexisting: the hypervisor allows the creation of several partitions; one of them could run z/VM with many virtual machines. At the same time, the workload manager could make sure that critical applications receive the correct amount of resources.

z/VM virtualization technology allows sharing the available hardware between the guest systems, leading to better utilization of the hardware. Deployment of new servers is very fast using cloning techniques because they are only logical definitions and can share the already available hardware resources. The consolidation of the different servers of a Web application can lead to a synergy effect of higher availability: with Linux for zSeries, not only the servers but also network and storage can be consolidated and simplified, reducing the complexity and making the environment more robust.

There are several cost advantages to consolidating distributed servers with Linux on zSeries:

- ▶ Distributed servers require many discrete servers: In even the smallest environment, there is a test server, a production server, and a backup server. With zSeries these can all be virtualized.
- ▶ Security can be enhanced by protecting the kernel and most of the binaries on a z/VM read-only minidisk.
- ▶ Upgrades to the server environment are easy to implement through cloning and sharing of environments.
- ▶ A company that is out of floor space will have to build a new data center if they continue along the current distributed path. This can be avoided with server consolidation.
- ▶ Large software charges from distributed software vendors can be avoided.
- ▶ Disaster recovery plans can be improved.
- ▶ Excess capacity on an existing mainframe can be utilized.

In most business environments, some good candidates for server consolidation with Linux with z/VM on zSeries are:

- ▶ File/print servers
- ▶ Database servers
- ▶ Web servers
- ▶ Network infrastructure
- ▶ E-mail

## 11.2.2 Storage implementation example

When it comes to storage, many companies are only making use of a percentage of their capacity. When storage is spread between servers, with no effective means of sharing capacity, reallocating or reconfiguring storage resources can be costly.

IBM Storage subsystems such as ESS, DS8000, DS6000 or DS4000 (formerly FASTT) families all provide Physical consolidation of the data, at the Block levels (disk level), with their capacity to provide logical drives (LUNs) from their available physical disks and share them between multiple servers across a storage network.

In fact, there are two established types of storage networks relevant to storage consolidation: the Network Attached Storage (NAS) and the Storage Area Networks (SANs). In simple terms, clients share files with NAS, and servers share storage with SANs. NAS is quick, easy, and inexpensive to implement because it uses the existing IP infrastructure. SANs are built upon dedicated, high-performance fiber channel networks specifically designed for large data transfers. Another option still in the early stage of its development is SCSI, which transfers storage commands (SCSI) over industry-standard TCP/IP to provide inexpensive storage networking. NAS and iSCSI solutions are both providing also a physical consolidation of the storage space, but now at a file level.

Announced in the fall of 2003, the SVC which is designed to help manage the complexity and costs of SAN-based storage, provides a Logical consolidation of the storage space also at Block level (block level aggregation). Based on virtualization technology, the SAN Volume Controller supports a virtualized pool of storage from the storage subsystems attached to a SAN. Its main benefits are:

- ▶ Centralized control for volume management: The SVC is designed to help IT administrators manage storage volumes on their SAN, helping to combine the capacity of multiple storage controllers — including storage controllers from other vendors — into a single resource, with a single view of the volumes.
- ▶ Avoidance of downtime for planned and unplanned outages, maintenance and backups: IT administrators can migrate storage from one device to another without taking the storage offline, and they can reallocate, scale, upgrade, and back up storage without disrupting applications.
- ▶ Improved resource utilization: The SAN Volume Controller is designed to help increase storage capacity and uptime as well as administrator productivity and efficiency, while leveraging existing storage investments through virtualization and centralization of management.
- ▶ A single, cost-effective set of advanced copy services: The SAN Volume Controller supports advanced copy services across all attached storage, regardless of the intelligence of the underlying controllers.

Now as NAS and iSCSI solution are both using IP as storage network, they are limited in performance compared to SAN solutions which are using Fibre Channel as the underlying network. SFS is intended to combine the benefit of file sharing across servers provided by the NAS architectures with the benefits of high performance data access provided by Storage Area Networks (SAN).

SFS is designed to provide network-based heterogeneous file system to support data sharing with policy-based file placement and file movement features in an open environment. SAN File System is also providing logical consolidation of the storage space such as SAN Volume Controller at the File level.

Its main benefits are as follows:

- ▶ It provides a common file system for UNIX, Windows, and Linux servers, with a single global namespace providing data sharing across servers through a shared file system.
- ▶ SAN File System clients are accessing to the data directly through SAN network (SAN File System is an asymmetrical virtualization solution).
- ▶ It improves productivity and reduce the pain for IT storage and server management staff by centralizing and simplifying management through policy-based storage placement and management automation.
- ▶ It can simplify and lower the cost of data backups through application server free backup, and by using built in file-based FlashCopy functions.

Now regarding the infrastructure management, IBM TotalStorage Productivity Center is an integrated storage infrastructure management solution that can help simplify and automate the management of devices, data, and storage networks.

IBM TotalStorage Productivity Center is comprised of the following four products:

- ▶ IBM TotalStorage Productivity Center for Fabric (formerly IBM Tivoli Storage Area Network Manager) brings all the information about SAN topology and configuration into a single place and is able to create, by correlation between different sources of information, topology mapping of the SANs that can be displayed graphically from a central management point. It can also provide other features such as automated device discovery, error detection fault isolation, SAN error predictor, zone control, real-time monitoring, and alerts that fully contribute to both reduce and simplify storage infrastructure management.
- ▶ IBM TotalStorage Productivity Center for Data (formerly IBM Tivoli Storage Resource Manager) is designed to provide a comprehensive Storage Resource Management (SRM) solution. By providing reports, monitoring and alerts, policy-based action, and file system capacity automation, it enables administrators to better identify, control, manage and predict storage usage, which simplifies IT infrastructure management.

- ▶ IBM TotalStorage Productivity Center for Disk (formerly IBM TotalStorage Multiple Device Manager — Performance Manager feature) is designed to centralize management from a single console of networked storage devices such as ESS and DS4000 product families or also SAN Volume Controller. Thus, TotalStorage Productivity Center for Disk helps reduce storage management complexity and costs, and also improves storage resource utilization.
- ▶ IBM TotalStorage Productivity Center for Replication (formerly IBM TotalStorage Multiple Device Manager — Replication Manager feature) is designed to simplify and improve the management of replication on network storage devices such as the IBM TotalStorage Enterprise Storage Server 'ESS, for which it provides copy services management.

### 11.2.3 BladeCenter implementation example

BladeCenter offers another opportunity for server and application consolidation by delivering integration, performance, manageability, resiliency, and investment protection in a blade architecture. At twice the density of today's 1U servers, BladeCenter enables a highly managed infrastructure that helps maximize resource productivity and minimize IT administration costs. Cost savings are provided by reducing footprint, cables, power, and server management. In addition, the ability to activate blades on demand provides both permanent and temporary capacity to meet the immediate needs of a business. The recent addition of the JS20 model blade based on dual PowerPC® 970 processors reinforces the concept started with the Intel Xeon blades.

## 11.3 Key products to start with


In summary, we can say that the following IBM products should be evaluated when needed to consolidate the infrastructure:

- ▶ IBM TotalStorage DS4000 storage subsystems (FAStT)
- ▶ IBM TotalStorage Enterprise Storage Server (ESS)
- ▶ IBM TotalStorage DS6000 or DS8000 Series storage subsystems
- ▶ SVC
- ▶ SFS
- ▶ IBM TotalStorage DS300 iSCSI RAID array
- ▶ IBM Ultrium Linear Tape Libraries (LTO)
- ▶ IBM TotalStorage Virtual Tape Server (VTS)
- ▶ IBM @server zSeries, iSeries, pSeries, xSeries
- ▶ IBM @server zSeries – zVM with Linux
- ▶ IBM @server xSeries with VMware

- ▶ IBM @server Blade Center







## How to optimize utilization and pool resources across a heterogeneous environment

IT organizations today have business applications that are multi-tiered and heterogeneous. Transactions and units of work supporting the business span multiple servers and systems. While faced with these types of environments, IT organizations within every industry are being challenged with reducing costs for IT resources as well as for reducing the amount of skilled labor to manage them, all of this while staying focused on ensuring that IT provides the resources needed to support the business.

To address the need to reduce costs, IT organizations must optimize their IT resources at higher utilization rates to achieve a better return on their investment. Doing this with today's heterogeneous infrastructures can be a challenge.

In addition to being able to optimize the resource, the needs of the business must also be considered. Additional workload cannot be randomly sent to servers that have lower utilization rates, since those servers may be positioned to handle a planned peak demand at a certain point in time. To guard against this happening, utilization optimization must be tied to the established policies regarding business application performance and service level objectives. The objective is to optimize in order to reduce costs, while at the same time maintaining the service levels required to support the needs of the business.

## 12.1 Vision

Ultimately, customers will be able to use technologies such as OGSA and Enterprise Workload Management to enable systems to dynamically participate in a heterogeneous, networked environment.

Management products continue to be enhanced to support workload management that spans hardware, middleware, and applications. The grid architecture and the workload management are two key components in this resource optimization issue.

### 12.1.1 Open Grid Services Architecture

Open Grid Services Architecture (OGSA) is an emerging standard that provides a standard technology base for managing and using distributed, heterogeneous resources. OGSA builds upon grid, autonomic, and Web Services technologies.

- ▶ Grid computing allows a customer to virtualize their resources, allowing the best possible utilization of resources as they are needed.
- ▶ Web Services provide standardized cross-platform integration of resources.
- ▶ Autonomic computing allows a customer to free up costly system management resources for strategic tasks while the infrastructure manages itself.

Continued development of grid standards through the Global Grid Forum will define standards built on OGSA, notably Grid Core Services, Grid Data Services, and Grid Program Execution Services.

As grid standards mature, an important capability will be the combination of orchestration capabilities with the application scheduling and workload management aspects of grid computing. Orchestration recognizes and dynamically responds to events based on business priorities, provisioning additional resources automatically. In effect, by sensing and responding to workload trends, orchestration provides an optimal infrastructure profile that can then be utilized by workload managers.

Both of these are focused on improving resource utilization, but take different approaches. Orchestration contributes to lower infrastructure costs and IT efficiency by optimizing the infrastructure to match workload profiles. Grid computing seeks to drive application workload onto idle capacity across the full breadth of available resources. Combining the two leads to a powerful effect on resource utilization.

## 12.1.2 Workload management

In order to address the need for workload management in a heterogeneous environment, IBM provides in the VE Suite for Servers the Enterprise Workload Manager (EWLM). EWLM interacts and interfaces with other components to manage a pool of servers. It will interface with provisioning and orchestration components to allocate additional resources as required.

At a storage point of view, by moving storage intelligence from the application servers or storage subsystems into the storage network (SAN), storage virtualization solutions allow a better polling or sharing of heterogeneous resources, at the disks level (block level) and at the file systems levels (file level).

## 12.2 How to get started today

In order to manage utilization and workloads across a pool of heterogeneous resources, workload management functionality is needed. This type of functionality is not something new to IBM.

The following sections describe some of the solutions that can be put in place already today in workload management.

### 12.2.1 The zSeries example

The S/390 architecture has, for years, provided this capability.

#### **z/OS Workload Manager**

*The Workload Manager* (WLM) component of z/OS does goal-oriented resource management by prioritizing all work depending on its importance, allowing different workloads to run within the same operating system and enabling them to share the same physical resources.

In addition, a physical S/390 machine can be divided into logical partitions running totally separate operating systems (servers), each with an assigned “LPAR weight,” representing the percentage of overall processing power assigned to all of the work in that partition. If workloads change and more processing power is needed in a particular partition, processing power can be shifted to the partition that needs it, as long as CPU cycles are available. If all partitions were at peak utilization, the operator of the system could change the LPAR weights manually. If the demand on the system was unpredictable and irregular (as in a Web server environment) and the system was highly utilized, the system would then need to be monitored at all times, to ensure that high priority workloads received the resources they needed.

## Intelligent Resource Director

*The Intelligent Resource Director* (IRD), which was introduced as part of the zSeries and z/OS, is designed to give an installation an enhanced ability to dynamically move resources to the most important work. This extends the concept of goal-oriented resource management by allowing the grouping of logical partitions that are resident on the same physical server into an LPAR cluster. This capability has been designed to give the z/OS Workload Manager the ability to manage resources, both processor and I/O, not just in one single image but across the entire cluster of logical partitions. But, this capability is all performed within a homogeneous environment.

### 12.2.2 Examples of LPAR

LPAR technologies are now available across the full @server product family. iSeries and pSeries servers both have enhancements for dynamic LPAR capabilities and xSeries servers use VMware ESX server to enable logical partitioning. Dynamic LPAR allows the dynamic movement of resources such as processors and memory. This capability is what is needed within a heterogeneous environment.

### 12.2.3 The grid benefit

Grid computing allows customers to integrate heterogeneous resources and improve asset optimization. Grid is enabled by a layer of middleware and services providing security, resource management, data management, and access to both static and dynamic information about the resources in the grid. Higher level services such as scheduling, brokering, accounting, and others are also an integral part of most grids.

As an example of how grid computing can be utilized, a financial organization might utilize a grid computing environment to perform risk analysis calculations that can be spread across many under-utilized servers. This provides significant efficiency and cost savings. By utilizing a grid to access additional resources, a parallelized application such as a financial analysis that used to take hours can now take minutes. Problems that were previously not practical to consider can now be solved.

IBM provides support for the OGSA standard through the IBM Grid Toolbox V3 (IGT3). IGT3 is an implementation of this standard, a new specification that the Globus Project played a key role in defining. Globus is using OGSA as their infrastructure for their GT3 base services (GT3). OGSA represents the reference open source standard implementation for the grid today:

- ▶ OGSA defines the blueprint of how a grid should appear, including the infrastructure; it defines the programming model of grid services; and it provides instructions on how to build a grid service by outlining the required components needed to build and deliver an enterprise-class grid solution.
- ▶ OGSA implementation is real. It is the middleware, the “Java 2 Platform” for grid services; and it defines how to build a grid service, outlining the mechanisms for creating, managing, and exchanging information for grid services.
- ▶ Web services are the foundation for grid services, which are the basis of OGSA, and therefore, GT3.

*The IBM Grid Toolbox V3 for Multiplatforms* is a commercial release of the Globus Toolkit Version 3.0 with IBM Value. It is a comprehensive, integrated toolkit for developing, deploying, and managing grid services. This product includes material developed by the Globus Project as well as a set of APIs and development tools to create and deploy new grid services and grid applications. It also includes a limited integrated hosting and development environment, capable of running grid services and sharing them with other grid participants, such as grid service providers and grid service consumers. It also provides a set of tools to manage and administer grid services and the grid hosting environment, including a Web-based interface, the Grid Services Manager. For more details on the Globus Project, see:

<http://www.globus.org/>

The IBM Grid Toolbox V3 provides customers with the following benefits:

- ▶ Quicker, easier development and deployment of grid services. A Software Developer’s kit (SDK) provides a collection of APIs. Tools allow the user to create a grid service from existing Java code or a Web Services Description Language (WSDL) interface, or to package services and related artifacts into a Grid Service Archive (GAR) for deployment into the runtime environment.
- ▶ Some key middleware technologies in addition to components harvested from the open source community. The infrastructure includes embedded versions of WAS-Express V5.0.2 for the run-time environment, IBM Cloudscape™ V5.1 database for storing and using working grid code, and OpenJMS for the JMS-based notification framework.
- ▶ Support for multiple platforms to meet the needs of heterogeneous environments. The toolbox supports AIX on pSeries, Linux on xSeries, OS/400 on iSeries, Linux on iSeries, and Linux on zSeries.

The grid services packaged with the IGT3 can be used either independently or collectively to develop useful grid applications and programming tools. Grid services are available for container management, logging, security information, data management, and program management. The toolbox includes several

services that are not provided in other open source implementations, such as the Common Management Model (CMM) and policies.

Additional benefits include:

- ▶ Wizard-based installation to streamline the install experience. This replaces the use of Grid Packaging Technology (GPT), used in many open source implementations, with native packaging and install technologies.
- ▶ Robust testing, including system/cross platform testing on the applicable IBM @server hardware.
- ▶ Additional services, including:
  - ServiceGroup, based on IBM XML Native Database technology.
  - A Web-based management application used to manage services within the runtime environment.
  - Common Policy Services: Common policy service manager (PSM) and policy service agent (PSA), which provide common components required for any policy-based application or solution.
  - Common Manageability Model Services to define and to model resource associations with service data. It provides mapping layer (abstraction) between existing systems management systems and technologies (CIM, SNMP). The initial release includes a JCA Adapter for CIM used to manage interactions with instrumentation.
- ▶ Documentation delivered via the IBM @server Information Center, including samples and tutorials to assist with education and understanding of the technologies and capabilities packaged within the product.

In the near future, both OGSA and Web Services technologies will converge to create a new open standard. The Web Services scope was firstly focused on Web applications to externalize pieces of existing applications and offer them to other applications or users. As the communication layer and the description of the offering was standardized (with SOAP and WSDL), the interoperability became easier. Java Web applications can now use other non-Java pieces of applications (such as Microsoft.NET), and vice versa.

The extension of this capability to resources other than applications can cover a wider scope of utilization. By offering resources such as computation or storage as a service, Web Services takes the grid strength, and the grid increases its flexibility. This new standard is named *Web Service Resource Framework* (WS-RF). WS-RF will be soon the most important communication protocol, and will be extensively used in the communication between the Infrastructure Management products.

## 12.2.4 Mixing partitioning and grid capabilities

With its capability to run many Linux applications simultaneously (including the IGT3) there is a great value in considering the use of the zSeries in a grid infrastructure. Some built-in features of the mainframe, such as Dynamic Server Provisioning, High Availability, Virtualization, and Server Consolidation are considered key attributes for grid solutions:

- ▶ Since the hardware itself, and its architecture, are designed for high availability (the Parallel Sysplex, for example), this means that all the components that are failure sensitive (both infrastructure and applications, grid scheduler and directory software, for example) can be placed onto the zSeries while facilitating high-speed access to mission-critical enterprise data and application resources.
- ▶ With both the LPAR and the z/VM-based support, the zSeries can be used for failover grid scenarios. If a Linux image fails, another image can be created quickly, in minutes, and a new physical box is not needed.
- ▶ Using the virtualization capability, a development and a testing environment can be created to develop and test the scalability of grid software. For example, one can create a central server to host the job submission interface, another server to host the Globus Monitoring and Discovery Services, a third server to host a broker, then create several work servers.
- ▶ One can use a grid node on Linux on zSeries to get access to other mainframe-based operating systems such as z/OS. The zSeries has a built-in high-speed network called HiperSockets. While the operating system just sees a network adapter, all the communication takes place in memory. With this capability one can connect very quickly and securely, since everything would be running in one box, to the database and transaction systems running on z/OS.
- ▶ Usually, grid tasks are run as separate processes within the same operating system, thus sharing resources controlled by the operating system. This situation might result in intentional or accidental exposure or corruption of the data of one task by another task. With the capability of running multiple concurrent virtual machines on a given zSeries, the level of isolation and security between grid tasks is unparalleled.

## 12.2.5 Infrastructure management tools

*IBM Tivoli Provisioning Manager (TPM)* provisions and configures servers, operating systems, middleware, applications, and network devices acting as routers, switches, firewalls, and load balancers. IBM Tivoli Provisioning Manager, through workflows, automates the manual provisioning and deployment process. Using pre-built “industry best practice” workflows to control and configure popular vendor products, or user-customized workflows to

implement environment or proprietary application-specific “best practices,” processes can be automated and executed in a consistent error-free manner.

*IBM Tivoli Intelligent ThinkDynamic Orchestrator (TIO)* extends the benefits of the IBM Tivoli Provisioning Manager. It orchestrates the activities necessary to automatically maintain server availability and meet required service levels. It provides the why, where, and when of a complete solution. By monitoring the applications under its control, IBM Tivoli Intelligent ThinkDynamic Orchestrator can sense degrading performance and determine what actions need to be taken. Because solutions are monitored closely, TIO can determine where (for which application) a resource is needed and instruct TPM to deploy a server, install the necessary software, and configure the network. This enables an application to maintain acceptable service levels.

With TPM or TIO, heterogeneous servers can be provisioned today. One can use TPM or TIO to install AIX on a pSeries server, Linux or Windows on an xSeries server, or Linux on a z/VM server.

This approach includes a description of the Enterprise Workload Manager components and their functions. The intention is to provide an understanding of how such a technology implementation would apply within a heterogeneous On Demand Operating Environment, and the benefits it brings.

*Enterprise Workload Manager (EWLM)* has been released in August, 2004. This product is part of the VE Suite for Servers.

Before we get into the details of EWLM, it is important that we understand why this technology is needed. Applications today are multi-tiered and each tier can be serviced by a different type of servers with a different operating system with different middleware. Having a transaction spread across many different tiers presents several problems and needs from an IT management perspective.

IT organizations need the ability to determine how the application is performing or whether it is achieving the service level objectives established for a business unit. In homogeneous environments, the ability to track a transaction end-to-end was fairly easy given how technologies evolved to support that need. If the performance was known, it was possible to determine whether or not service levels were being met.

Today's environments need to be *managed end-to-end* from the standpoint of performance as well. The mainframe-based WLM type of workload management functionality — which balanced system resources and ensured that the proper ones were applied to the most important business applications in that environment — is also needed within the heterogeneous environments that are common to businesses today. EWLM is the result of this adaptation work.



The second need is related to IT organizations working to reduce costs by increasing the utilization levels of servers. Technology will be needed to identify the utilization levels for a given server before it can be determined if any adjustments should be made to workload on that server. Enterprise Workload Management is the technology that will provide this type of information within the On Demand Operating Environment.

## Enterprise Workload Manager focus

Enterprise Workload Management technology provides the functionality to manage the heterogeneous servers that are defined within what is called an Enterprise Workload Management domain. Each Enterprise Workload Management domain can have hundreds or even thousands of server resources within it.

Enterprise Workload Manager basically consists of four pieces:

- ▶ **Domain policy:** These are the policy-based service level objectives for the workload that will run on the servers within the EWLM domain. The objectives are defined by establishing domain policies and service policies. These definitions are entered using the administrative user interface and are sent to the domain manager where they are stored in XML format.

**Note:** The policy definition is very close to the zWLM one. Terminology and components are the same. Hence, zSeries skilled people won't have any problem to extend their zWLM experience to EWLM. For others, this is a new theoretical aspect which should first be worked on.

- ▶ **Domain manager:** There is one physical domain manager for each Enterprise Workload Management domain of servers that are being managed. It is responsible for storing the domain and service policies. In addition, the domain manager dynamically manages the server topology within the EWLM domain and maintains it along with the state of each of the servers. It also handles communications with the administrative user interface (a Web based GUI) as well as with all servers in the EWLM domain.
- ▶ **Managed server component:** This is code that runs on each server within the ELWM domain being managed. Each server contains an implementation of the code, which is based on the Application Response Method (ARM) standard from the Open Group<sup>1</sup>, either delivered with the operating system or installed along with EWLM, depending upon the platform. This EWLM ARM implementation interfaces with the operating system as well as with instrumented software running on the server. Information is provided which allows the tracking of workload that is flowing through the servers in the EWLM domain. Information that is gathered is sent to the domain manager to

<sup>1</sup> See more information at the following link: <http://www.opengroup.org/tech/management/arm>

be summarized and analyzed. This information gives insight into how the server is running from a utilization standpoint as well as how well the pieces of the workload are being serviced. The information is then available to be accessed from the administrative user interface.

- **Administrative user interface:** This is the point from which the EWLM domain is controlled. The domain policy and service policies can be entered from this interface. The interface is also the tool used to control when domain policies and service policies are to be implemented. This interface can be used to instruct the domain manager to implement certain domain policies and service policies for the environment. The domain manager then communicates with the managed server components running on each of the servers.

By being able to know the utilization of a server and the service level objective for the workload running on that server, it is possible to determine if the service level objectives are being met for the workload as well as whether the server can be utilized at a higher rate while not jeopardizing the SLA for the workload.

Implementation and exploitation of a technology such as EWLM along with the IBM Tivoli Intelligent ThinkDynamic Orchestrator and IBM Provisioning Manager can provide the following benefits:

- Allow performance measurement and problem isolation for groups of transactions running within multi-tiered business applications in heterogeneous environments.
- Reduce costs and improve the return on investment associated with servers by being able to increase the utilization of pooled servers while maintaining service level objectives.
- Reduce the number of skilled personnel and time required to manage the performance of the multi-tiered server environments by automating the orchestration and provisioning of systems to support service level objectives.

**Note:** For more information about EWLM, you can read a redbook dedicated to the topic: *IBM Enterprise Workload Management*, SG24-6350.

## 12.2.6 Storage focus

Related to storage now, IBM Totalstorage virtualization products, with their implementation into the SAN itself (SAN based solutions), by their central management architecture and by the native feature they provide can fully help to pool heterogeneous network-attached storage resources. Both products provide a *logical* consolidation of the storage space by opposite to the *physical* consolidation which is mainly implemented in the storage subsystems.

## IBM SAN Volume Controller (SVC)

IBM SAN Volume Controller (SVC), is designed to provide virtualization, aggregation and *logical* consolidation of the storage space at the block-level (disk level). It integrates also a Volume Manager for disks across the SAN. In simpler terms, this means that the SAN Volume Controller manages a number of back-end storage controllers (storage subsystems) and maps the physical storage within those controllers to logical disk images that can be seen by application servers in the SAN as LUNs, exactly if they had been defined directly on the storage subsystems themselves.

To do that, physical storage space has to be logically grouped, by the administrator, in storage pools named Managed Disk Groups under SVC terminology and from which can be then created the logical disk images called Virtual Disks under SVC terminology. This is really how is implemented the physical disk abstraction (or virtualization) by the SAN Volume Controller. The product allows creation of Managed Disk Groups from single such as multiple storage subsystems. It allows also the Virtual Disks to be moved non-disruptively, between Managed Disks Groups and without any constraint regarding the type or the vendor of the storage subsystems used in these pools. With this SVC major data migration feature, it is really possible to move data between heterogeneous storage subsystems (IBM or non-IBM) without any impact and down-time for the application that are using them.

## IBM SAN Filesystem (SFS)

IBM SAN Filesystem (SFS), now is designed to provide virtualization, aggregation and logical consolidation of the storage space at the file-level (file system level). It provides a high performance shared SAN file system, accessible under a Globale Name Space, with all necessary consistency, integrity and distributed lock mechanisms required by this kind of product. By moving this file system into the storage network itself, SAN FS provide a high-performance data sharing at the file system level for heterogeneous servers running different operating systems, Unix or Window. By allowing servers to share and pool more efficiently storage resources, SAN FS limits the number of duplicate files and reduces the overall amount of storage space required in the infrastructure.

SAN File System is also designed to manage available storage space as a unique “storage reservoir” (for information, the original SAN FS project name was “Storage Tank”); this global space has to be subdivided logically by the administrator in several Storage Pools regarding characteristics of the underlying storage subsystems LUNs. Follows some examples of the possible criteria that could be chosen to define these storage pools:

- ▶ Device capabilities (RAID level)
- ▶ Performance
- ▶ Availability

- ▶ Location: secure or unsecured
- ▶ Business owners
- ▶ Application types

Once this storage repartition has been performed, SAN File System policy based File Placement mechanism could be used to control physical placement of files, as they are created in the SAN FS Global name space. This placement in the different available Storage Pools is managed by a set of rules (grouped in policies) based on many possible files attributes that can be:

- ▶ File name
- ▶ File extension
- ▶ File owner
- ▶ Creation date
- ▶ And so on...

For example, it would be possible to have a critical projects directories stored in a high performance and available storage pool, users personal directories stored in an other storage pool and all files with extensions \*.mp3 and \*.avi stored in a low cost storage pool. This is the case in the Video on Demand (VoD) sector.

The latest release of SAN File System integrates also some new File movement and Life Cycle Management features.

- ▶ The “File movement” feature allows the administrator to move files non-disruptively between storage pools. The benefits of this feature is to be able to change the class of storage to match the business value of a file and effectively change the pool-dependent settings used by a file to improve performance. It can also be used to redistribute a file to stripe across newly added volumes in a storage pool, or to correct the outdated or unintended effect of a file placement policy rule.
- ▶ The second feature allows automatization of file movement using file management policy-based mechanism. This feature provides life cycle management capabilities to SAN File System. A file management policy is a set of rules, which specify conditions for either automatically moving files from one storage pool to another, or for deleting them entirely from the storage pool. This life cycle management feature will also permit a better utilization of the storage space by automatically move (or delete) files to appropriate storage pool depending on attributes that can be size and/or file age (number of days since the file was last accessed). For example, we can decide with this feature to have all files not modified since six months to be automatically moved to a slowed but less expansive storage pool and freeing then space in faster storage pools for current projects files.

## 12.3 Key products to start with

The challenge of optimizing the utilization of server and storage resources across multi-tiered implementations in heterogeneous environment can be addressed as technologies evolve. Ultimately customers will be able to use technologies such as EWLM and others that will allow systems management control from a centralized console. These technologies will enable systems to dynamically participate in the heterogeneous, networked environments of dynamic on demand companies.

The following IBM products should be evaluated for their use in developing a solution using this approach:

- ▶ IBM TotalStorage SAN FileSystem
- ▶ IBM TotalStorage SAN Volume Controller
- ▶ zSeries with Intelligent Resource Director
- ▶ z/OS Workload Manager (z/OS WLM)
- ▶ IBM Enterprise Workload Manager (EWLM)
- ▶ IBM Grid Toolbox





## How to provision system resources in order to meet business demands

Provisioning in an On Demand Operating Environment is the capability to automatically deploy and dynamically optimize operational resources in response to business objectives in a heterogeneous environment. There are many types of resources that may be provisioned: servers, storage, identities, applications, networks, and more.

When platform provisioning technologies are combined with other provisioning technologies and enhanced by orchestration, the On Demand Operating Environment will have the ability to make the most informed decisions about provisioning hardware, software, applications, and so on to optimize the IT infrastructure.

Businesses need to be able to address the following requirements:

- ▶ Quickly deploy new N-tier applications and all of the resources that support them to stay responsive and competitive.
- ▶ Support a heterogeneous environment, with multiple hardware platforms, operating systems, and software and middleware components.

- ▶ Reduce administrative costs associated with deployments through the use of automation.
- ▶ Increase system utilization by quickly and reliably re-provisioning systems to meet immediate business requirements, including service level objectives.
- ▶ Make the underlying hardware and OS platforms as transparent to the applications, business processes, and users as possible.

## 13.1 Vision

As technologies and products evolve, we might expect to see the same flexibility on all hardware platforms, as far as subprocessor allocation goes, and with respect to the number of independent operating systems hosted.

The workload manager that today runs within the operating system may find its way to the hypervisor to be able to balance resource across partitions, and eventually across different systems and architectures.

The software component of provisioning solutions should also adapt to these changes, as it becomes able to interface with the global workload manager and improve its capability to define more complex enterprise policies.

Another important aspect of the On Demand Operating Environment related to provisioning is billing and metering capabilities. The ability to automatically measure the resources consumed by task, user, or department is critical to ease the billing process.

## 13.2 How to get started today

Refer to Chapter 16, “How to provision system resources according to business demands” on page 215, for a more complete discussion of this approach.





## How to monitor end-to-end applications, their topology, and their resources

One important area that contributes to infrastructure management is the capability to understand the availability and performance of user transactions or applications. Data collected on application performance and availability can be used to determine what changes may be necessary within the IT infrastructure to meet business objectives for individual applications as well as the infrastructure as a whole.

One of the key metrics required for overall systems management is response times for the overall transaction or business process as well as the application components that make up the transaction. Because of the importance of this data (and the associated data, such as availability and throughput), IBM has chosen to implement the OpenGroup Application Response Measurement (or ARM) standard throughout our hardware and middleware portfolios. Implementation of this standard enables the real time measurement of transactions, transaction components, and underlying resource usage (hardware and OS) associated with the execution of the application.

## 14.1 Vision

To insure a coordinated and synergistic delivery of function, the automation blueprint (Figure 14-1) is used to define the product roles across the automation layer of the On-Demand blueprint. The automation blueprint coordinates the definition, delivery and effectiveness measures of the different autonomic actions. Mapping the capabilities today of TMTP and EWLM to the Automation Blueprint shows how they will leverage and support the overall On-Demand message.

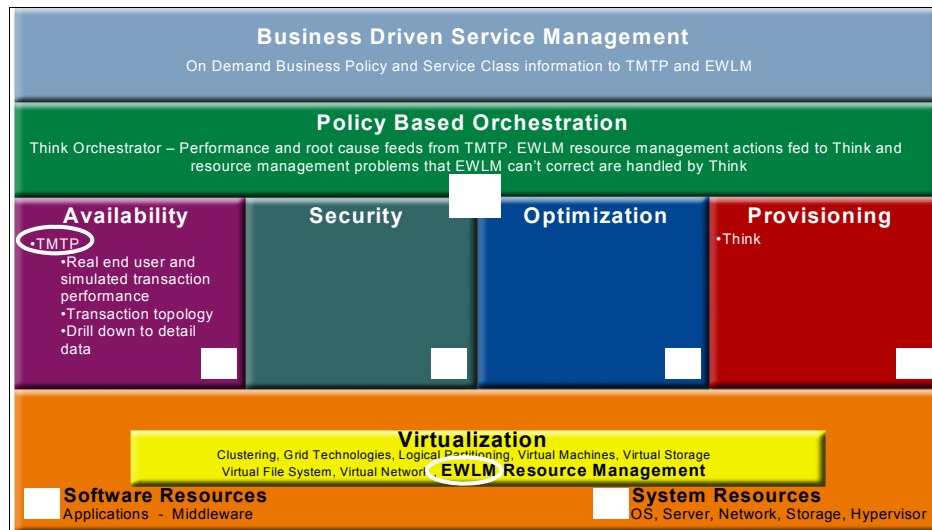


Figure 14-1 Automation Blueprint

In Figure 14-1, starting at the top, the definition of the business policy will be handled by the Business Driven Service Management layer. The execution of the autonomic cures to enforce the business policies need to be coordinated by the Think Orchestrator to insure conflicting actions are not taken. TMTP provides the end user perspective and the operational reporting and drill down views as well as providing data feeds to Tivoli Service Level Advisor (TSLA), the Tivoli Data Warehouse (TDW) and the Think Orchestrator. EWLM will implement resource management and traffic shaping actions and will integrate with Tivoli Orchestrator product to recommend provisioning actions. A standard based Service Provider Interface (SPI) will be provided for the ARM data collection so that other IBM or third party products can leverage ARM data being generated by the applications and middleware.

The combination of TMTP, EWLM and the Tivoli Think product set provides a comprehensive solution for maintaining service levels in an On Demand world. TMTP provides an early warning of performance problems and identifies the root cause. EWLM and Think can quickly take action to maintain service levels. EWLM can optimize the existing resource to support business priorities and Think can quickly add additional resources as required.

## **14.2 How to get started today**

Today, you can start addressing some key elements of the overall picture that will start the transformation of your IT infrastructure into an On Demand Business.

Because IBM has standardized on using the ARM interfaces, tools that exist in the operating environment can all leverage the same set of base measurements. Two of the products that will utilize ARM information are Tivoli Monitoring for Transaction Performance (TMTP) and Enterprise Workload Manager (EWLM).

### **14.2.1 IBM Tivoli Monitoring for Transaction Performance**

IBM Tivoli Monitoring for Transaction Performance (TMTP) is an operational product that is used to analyze performance and availability from an end users perspective and to assist customers in determining the root cause of performance problems. There are two basic ways of understanding the end user view of transactions. Active monitoring, which records and plays back simulated transactions and passive monitoring where data on actual end user traffic is collected.

The TMTP product supports both types of monitoring. While both active and passive monitoring provides an understanding of transaction performance from the end users perspective they do not provide much detail on the root cause of problems. A simulated transaction can tell you what step is failing or performing poorly, but in today's complex environments that could still leave a large number of possible problem areas to investigate. What is needed is a way to follow the transaction execution path and see exactly where problems are occurring. This transaction decomposition function is where TMTP leverages ARM.

Transaction decomposition involves following the transaction as it flows through the environment and understanding the performance in detail across the systems on which it runs. TMTP provides intuitive and operator friendly topology maps and interactive graphs to quickly and easily display the root cause of performance problems. The transaction topology maps are dynamically discovered based on customer specified URLs or J2EE components they wish to be monitored.

Performance thresholds for each monitored component can be automatically set to enable a quick time to value for the customer. From the high level transaction view, an operator can quickly drill down to individual JAVA methods or SQL calls to determine the root cause of a problem, eliminating lengthy problem determination cycles. This transaction decomposition function can be tied to the active and passive monitoring components or used as a standalone function. TMTP also ties together the transactional view and the resource view by providing a direct link, in context, to the IBM Tivoli Monitoring (ITM) Health Console which provides resource level views of problem components.

The information from TMTP will provide an end user performance and availability perspective that can be used when establishing an On-Demand business policy. The business policy will define the performance and availability goals of the transactions, which can be grouped into service classes and assigned a relative importance based on business goals. The information from TMTP and the business policy can then be used to drive autonomic actions, such as workload management or provisioning, to ensure the business policy goals are satisfied.

## **14.2.2 Enterprise Workload Manager**

Enterprise Workload Manager for Resource Optimization (EWLM) will leverage the same base of ARM information to implement autonomic cures from an overall system perspective. The goal of EWLM is to provide policy-based, goal-oriented resource performance management. The implementation begins with a business policy definition as defined above. The relationships are then quite simple: Satisfy the goals of the most important service classes, and then allow the remaining work to absorb whatever resources are left.

This support consists of algorithms that absorb the collected performance statistics, learning which classes of service are satisfying their goals, and which classes are not. The logic also tracks relevant resource relationships, learning where the application of additional resources (for example, CPU, memory, or network bandwidth) can help, and where it may not. When the goal for a service class is not being satisfied, EWLM working with system and network resource managers, attempts to fix the problem while insuring that the changes made do not cause more important workloads to exceed their performance goals.

The EWLM approach builds upon the z/OS WLM support, applying many of the same techniques to other platforms and to other situations. EWLM learns about relationships between processes, between applications, and between server and network resources. Using knowledge about the flow of work throughout the configuration, it can decide on which servers help is possible, and on which servers any attempts to help would likely be futile. This support can also extend beyond operating systems to network traffic shaping and load balancing using the WebSphere Edge Server.

You can refer to Chapter 19., “How to monitor using EWLM” on page 295 for an EWLM sample scenario on how to set up the EWLM infrastructure and a sample demo on how EWLM can be effective in monitoring your enterprise applications.

## 14.3 Key products to start with

The following IBM products should be evaluated for their use in developing a solution using this approach:

- ▶ IBM Virtualization Engine Suite (EWLM Component)
- ▶ Tivoli Monitoring for Transaction Performance (TMTP)





## Part 3

# Infrastructure management: Detailed scenarios

The previous part of this book has provided an overview of different approaches or “how to’s” that an enterprise might use to start implementing an On Demand Operating Environment today. These how to’s cover various aspects of the infrastructure management components of the On Demand Operating Environment. However, other approaches are also possible. We chose these as representative of the types of issues many enterprises may be facing today.

For each approach, we provided an overview of the business requirement, a vision of the type of functions that will be required and available in the future, and a short description of how one can get started today with products and capabilities that are already available.

In this part of the book, we describe in detail five example scenarios of how an enterprise might plan for and deploy a solution to put in place today an On Demand Operating Environment while keeping in mind and preparing for the longer term vision. These scenarios are based on the vision and the components described in Part 2, “How to’s for managing the Infrastructure” on page 67.

There is no specific significance to the approaches we chose to detail in the following chapters. Rather, we just wanted to provide examples of how one might start the detailed planning and deployment of an On Demand Operating Environment based on any of these approaches.


Each scenario wants to address a specific area of the on demand world. However all the areas of the On Demand Business are not covered here. The objective was not to cover in this part each technology or product referenced previously. We aimed to give a practical approach on very common issues, so that the readers, through these scenarios, can better understand how an infrastructure can evolve towards an on demand infrastructure.

The chosen examples address the following infrastructure management area:

- ▶ Chapter 15, “How to secure access and control of information, resources, and applications” on page 167, addresses a security issue.
- ▶ Chapter 16, “How to provision system resources according to business demands” on page 215, addresses the problem of under utilized resources, by provisioning.
- ▶ Chapter 17, “How to balance workloads in the network” on page 255, address how to make the most of the resources, by balancing the workload.
- ▶ Chapter 18, “How to consolidate, simplify, and optimize the storage IT infrastructure” on page 271, addresses the consolidation problem, with a storage infrastructure example.
- ▶ Chapter 19, “How to monitor using EWLM” on page 295, addresses the end to end monitoring management.

**Note:** The first and previous version of this book, SG24-6634-00, has already addressed two examples, as written in Chapter 15 and Chapter 16. In this second edition of the book, SG25-6634-01, this part of the book has been extended with the examples in Chapter 17, Chapter 18, and Chapter 19.





## **How to secure access and control of information, resources, and applications**

This chapter focuses on the need for automated security to support an On Demand Business. As companies move forward in their strategies to become more responsive and resilient, it becomes increasingly important to secure access to and control of information, resources, and applications.

Automated security will become a critical element in each company's IT infrastructure. Automation will allow an IT organization to react with speed as their business reacts to marketplace changes, without introducing risk to the business.

This chapter identifies business needs related to security in an On Demand Operating Environment and identifies the products available today from IBM that can be used together in a solution to address these business needs. It also presents a practical scenario to highlight the concepts as they may apply to real-world enterprises.

## 15.1 Introduction

Information in data stores such as databases, transactional systems, file systems, and even Web Services is becoming increasingly accessible within enterprises as well as from the external world. The same applies to enterprise operational entities such as applications, business process workflow, operating system processes, and traditional batch jobs. The IT challenge is to secure and control access to information and operational entities automatically in real time. This requires synchronization of identity across the infrastructure as well as a consistent set of policies regarding access control. The goal is to provide controlled, yet pervasive access to information and applications as it is needed based on the demands of the business. The objective is to allow for the dynamic mapping of identities to data sources in real time, thus minimizing any delay in making data available to authorized users.

In practical terms, securing the On Demand Operating Environment requires a distributed security mechanism for authentication and access control within and across organizational boundaries. To provide consistent access, the security architecture for the On Demand Operating Environment supports, integrates, and unifies popular security models, mechanisms, protocols, platforms, and technologies to enable a variety of systems to interoperate securely.

The security of the On Demand Operating Environment is built upon the Web Services security specifications. For more details, see the Web Services Security Roadmap at:

<http://www.ibm.com/developerworks/security/library/ws-secmap/>

Specific bindings for security will provide protocol-specific details and security functions such as confidentiality, integrity, and authentication. For instance, transport bindings over a Secure Socket Layer (SSL) connection provide a secure connection between two end points: WS-Security is used to secure SOAP messages and WS-SecureConversation establishes a security context and derives session keys.

As businesses become more dynamic, so do their security, configuration, and user provisioning requirements. In a dynamically provisioned environment, new users may need to be added or deleted. Group access control lists may need changes. Such changes need to be reflected within the infrastructure, and synchronized, so that the security policies can be effective immediately. Similarly, policy changes such as the granting or revocation of access rights, and changes in membership and profiles that affect access rights, should be capable of being changed dynamically, enabling runtime decisions in real time to allow pervasive yet selective and secure access to resources. Web Services-based specifications such as WS-Policy and WS-Security Policy provide the basis to unify already existing policy models.

## 15.2 General strategy

The components needed to implement this approach have to respect the general attributes of an On Demand Operating Environment, specifically that it be:

- ▶ **Self-managing:** When any modification is brought to the environment (for example, employee hiring, job move, end of contract, new application being rolled out, new systems being deployed), and is accepted at the company policy level, the security aspect of the change must be reflected, and enforced with limited human intervention.
- ▶ **Scalable:** The solution that is deployed must be able to adapt itself to an increasing number of parameters (meaning number of users, number of groups, type of information stored for any objects, and so on), and also able to respond to an increasing number of requests for user or data authentication.
- ▶ **Resilient:** The security components used to implement this approach are a mandatory part of the complete solution. If security is faulty, then users won't be able to access critical information or applications and the business will suffer. For that reason, these components need to be resilient, meaning that the software and hardware parts must be able to react to an interruption of service by relocating themselves in a different portion of the IT environment.
- ▶ **Economical:** Any solution today needs to be economical, and a security solution is no different. The cost of the hardware and software is not the only factor; the overall cost of providing security administration needs to be understood. A centralized solution that can provide a consistent application of the enterprise's security policies in an automated fashion helps reduce the administrative overhead of keeping an IT environment secure.
- ▶ **Open standards-based:** The only way that a centralized solution can provide consistent enforcement of policies related to a large number of platforms, operating systems, applications, and data stores is by building on and using open standards. New components can be added to an environment where security is based on widely accepted open standards without the need to modify or adapt the security model.

In the security area, many open standards have already been adopted. Among them, we find Secure Socket layer (SSL), Transport Layer Security (TLS), Private Key Infrastructure for X.509 V3 (PKIX), Lightweight Directory Access Protocol (LDAP). In addition, Web Services Security standards are evolving and are being adopted as well.

## 15.3 Solution components

The four main capabilities that should be addressed by a security solution are:

- ▶ Identity management
- ▶ Privacy management
- ▶ Security management console
- ▶ Data protection

These capabilities are described in detail in the following sections.

### 15.3.1 Identity management

In an On Demand Operating Environment, swift and easy access to information is an important driver of business success. Making this information available to distributed business users has become a challenge to companies because the Internet-enabled rise of e-business has dramatically increased the scope and number of individuals requiring access. For instance, data may need to be accessed by an organization's customers, employees, partners, and suppliers. As the IT user base grows, so too must the number of applications employed, making security and user management complex and expensive. If each application has its own directory of identity data, requiring manual administration, the cost and efficiency of providing secure access to data to the right people and applications becomes an immense challenge.

The security solution that is adopted must ease the mapping of the enterprise security policy to user account creation. That is, each new user has a role and based on that role, they should have rights covering information access and resource access. Automating the granting of the appropriate access across all resources in an enterprise is a key component of an on demand solution.

Another important and often time-consuming task in security is password management. Generally, a user has different passwords on different systems, with different rule sets (number of characters, password lifetime, and so on) which make it difficult for a user to track all required passwords. Users then tend to simplify their passwords and often right them down, thus causing security exposures. When users forget their passwords, they often have to contact their security administrator, who resets them. The ability to have a single sign-on and automated processes for changing or resetting passwords if they are forgotten can greatly reduce the administrative overhead associated with identity management.

As the number of applications and resources that a particular user or group of users should be able to access increases, a solution must also support delegation of authority. Of course, this in turn requires appropriate logging and auditing facilities.

### 15.3.2 Privacy management

More and more enterprises share information with their employees, customers, partners, and suppliers. As the scope and number of these business users grows, the number of confidential transactions increases as well. Fast, easy access to information and transactions in a distributed environment has become an important component of how competitive businesses succeed. However, along with increased information access, comes the added requirement of managing the use of that data. Businesses must put in place policies and procedures to maintain control of information access and ensure they meet their own as well as government and regulatory policies related to privacy. If done manually, this can be an expensive and complex task, as the following list illustrates:

- ▶ **Cost:** Manual procedures to manage the privacy preferences of business constituencies can be time-consuming and costly.
- ▶ **Security:** Written policies and manual procedures managing the privacy preferences of customers, employees, and business partners increases the risk of inappropriately disclosing personal and confidential information.
- ▶ **Efficiency:** Manually publishing changes to existing privacy policies or adding new policies to applications and IT systems is time-consuming and ties up IT resources in non-revenue-generating activities.
- ▶ **Time:** Generating reports for privacy inquiries, audits, and regulatory compliance can be time consuming.

An On Demand Operating Environment provides automated solutions to help address this critical area.

### 15.3.3 Security management console

Having a security management console is critical in today's environment. Not only does it simplify user account administration, it also helps an enterprise react more efficiently in case of attacks.

Businesses face internal and external security threats from a multitude of fronts; viruses, unauthorized access, denial-of-service attacks and other forms of intrusions target applications, networks, hosting infrastructure, servers, and desktops. The high volume of security events that are generated can make it difficult to quickly identify and respond to real security threats.

A scalable solution must centrally manage security incidents and vulnerabilities from a single security console to provide an overall view of the security of the infrastructure. The solution must be powerful enough to control access to data across heterogeneous systems and applications.

### 15.3.4 Data protection

The data protection element has multiple aspects. It must include protection of data in case of loss (hardware incident or human error) through a backup and recovery solution. The solution must be able to protect data wherever it is, on local workstations or servers, in files or databases, on SANs or network-attached storage.

Data protection must also cover controlling access to the data by the users, and making sure that data access rights are enforced.

The security infrastructure should also allow transparency of the user's location. More and more users are not connected from within the company but may connect from outside, working from home, from a customer/partner site, or from a hotel while on a business trip.

## 15.4 Scenario

So far we have discussed the strategy and basic components of the automation approach for secure access and control of information, resources, and applications. Now, we will apply those theories and guidelines to a simple On Demand Operating Environment scenario for a hypothetical organization with a typical set of requirements.

### 15.4.1 Business context

ITSO-Electronics.com is a wholly owned subsidiary of a major brokerage company, Medvin, Lasser & Jenkins (ML&J). ML&J's online presence has, to this point, been limited, consisting mainly of informational Web content. Online trading has not been a priority. The clientele has traditionally been major accounts with assets greater than \$5 million, and transactions are almost exclusively done via direct contact with a broker. While the company, a privately held corporation, has maintained solid profitability over the past several years, largely due to a stable client base, the company's growth has stagnated, remaining at approximately the same revenue levels for the last several years.

Market trends have forced a rethinking of ML&J's approach to business. The individual investor community has increased substantially in recent years, and the company has not shared in that growth. Consequently, the company's market share has eroded. Also, the rise of online trading has begun to affect a portion of ML&J's client base. In the last year, there has been a net outflow of investment funds cutting across approximately ten percent of all client accounts. Research has shown that 95 percent of these outflows are being redirected to online brokerages. This trend, if it continues, threatens to affect the long-term viability of the business.

An online component to complement ML&J operations has been judged a necessity. ITSO-Electronics.com was started with assets recently acquired from a failed Internet startup, and additional capital has been provided to fund completion of the company, which recently began full production operation ramp-up. ITSO-Electronics services the online trading requirements of ML&J's current clients, while focusing on developing additional clients who are primarily online traders with trading capital in excess of \$250,000.

**Attention:** As of May 2005, the time of writing this section, our *fictitious* domain name ITSO-Electronics.com was not reserved by anyone.

## 15.4.2 Current environment

The ITSO-Electronics.com concerns for becoming an integral part of the ML & J IT infrastructure fall into three major categories:

- ▶ Data centers
- ▶ Network
- ▶ Operational plans

A closer look at these individual aspects is provided in the following sections.

### Data centers

ITSO-Electronics.com has two major data centers. One is located in San Diego, California, and the other is in Savannah, Georgia. At this time, all Internet application access and key internal application access is provided through the San Diego center, in which the company's IT Operations (OPS) group is based. The Savannah center is currently supporting a few other internally used applications and houses the company's IT Architecture, Development, and Deployment Support (ADDS) business unit.

While ITSO-Electronics.com did consider hosting its Web servers through a third-party provider, it was decided that all subsidiaries deploy its servers in-house. However, they have not ruled out migrating certain Web operations to a hosting provider in the future. This could bring additional data centers into play.

### Network

The data centers are connected by redundant T3 (45mbps) access. At this time, Internet connectivity is provided through the San Diego center, with multiple T3 lines from three different providers. Figure 15-1 shows the national ITSO-Electronics.com network.

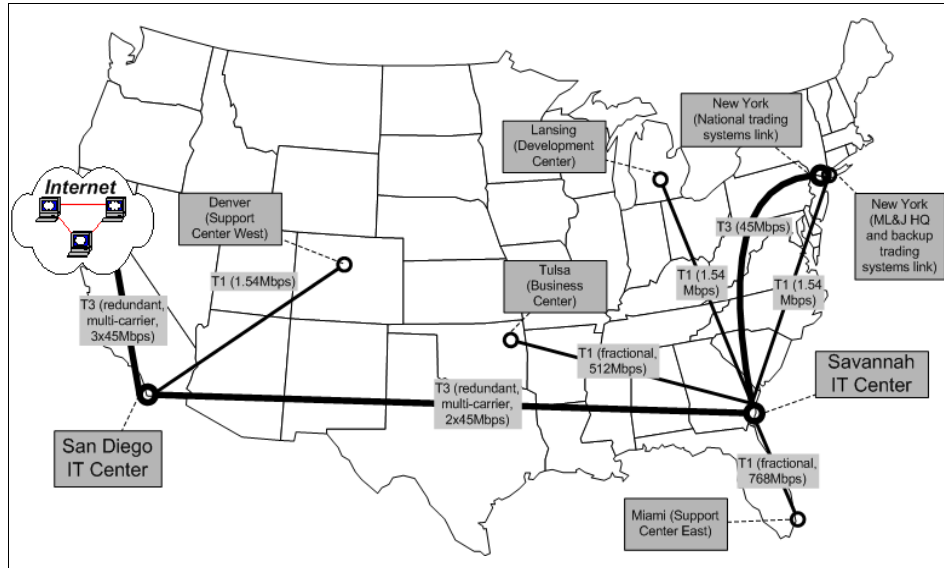


Figure 15-1 ITSO-Electronics.com data network

Within the San Diego center, all Internet access is channeled through Web servers residing in a demilitarized zone (DMZ). These Web servers provide front-end application logic, including presentation services. Back-end application logic is hosted on systems residing behind the DMZ in an internal production network.

The Savannah center has no direct Internet access. It has a production network for internal application systems.

In addition to the specific network capabilities at each of the sites, there is also a general company intranet shared across all corporate locations. This network is not considered secure, and is not authorized for hosting production systems.

## Operational plans

Early plans are in the development stage for future expansion of Internet operations into the Savannah center to provide for redundant access capability, with load-balancing for customers on the US East and West coasts. At this time, there is no requirement to actually support this. However, the ITSO-Electronics.com chief architect wants to be certain that the security solution they deploy is capable of meeting such a requirement.



### 15.4.3 Business objectives

The CIO has provided input on the business drivers for the targeted solution:

- ▶ Provide an enabler for consistent application of security policy across the business. The business cannot afford to create multiple, competing security infrastructures.
- ▶ Assure client confidence by offering a flexible, yet perceptively secure solution. It is essential that the security system not get in the way, while at the same time protecting client information and assuring that financial transactions are conducted securely.
- ▶ Competitively position the business to react quickly in deploying secure premium services and content. Quickly deploying value-add capabilities is important to gaining and maintaining market share.
- ▶ Allow for the integration of special premium applications capabilities to Medvin Lasser's "Select" clients. Medvin Lasser is very focused on maintaining their existing high-income client base by providing them with special capabilities that are not available through any other online service. For example, additional bond management capabilities within the portfolio management application are being developed specifically for these clients.
- ▶ Provide for expansion of services with minimal incremental investment. It is essential that, once in place, the security solution grow with the company. It is unacceptable to require extensive and continuing re-engineering efforts for the security infrastructure as the company expands its operations.
- ▶ Meet applicable US Securities and Exchange Commission (SEC) requirements. There are certain legal requirements for assurance that client assets and transactions are handled properly. The security infrastructure should be supportive of these requirements.

### 15.4.4 Technical objectives

ITSO-Electronics.com has been deployed as a Web-based online trading system with capabilities similar to those found at other online trading sites. This software is composed of a number of underlying applications, all of which perform functions based upon each user's privileges. For example, only users who have paid for "Level II" quotes may access that application.

In concert with the ongoing application development activities, the company has been examining alternatives for providing secure access to their Web site. Originally, a “master” application was developed, that provided a single access point for providing user authentication and authorization to use the underlying capabilities of the operating system.

Following initial deployment, additional requirements became apparent. It became clear that the level of effort required to fully address all functional requirements was cost-prohibitive. The tie-in to the operating system security mechanisms began to limit certain deployment options. The CIO felt that this approach was locking them in architecturally to an in-house solution that would require long-term maintenance and support services.

After examining marketplace alternatives in a proof-of-concept (POC) setting, a decision was made to move to the On Demand Operating Environment automation approach to secure access and control of information, resources, and applications.

The company wants to transition its user base from the in-house Web security system to one based on Tivoli Access Manager over the next several months. They initially wish to deploy adequate capacity to address their anticipated load over the six months, and then incrementally add more as needed.

After several discussions and review with one of the founders of ITSO-Electronics.com, Mr. Jenkins, the planning team identified the issues that ITSO-Electronics.com is facing today. These issues are explained in the following sections.

## **Too many inefficient processes**

First of all, there are too many inefficient processes within ITSO-Electronics.com.

- ▶ **Problem 1:** ITSO-Electronics.com uses slow and inconsistent processes to create user accounts and provision user access rights, as illustrated in Figure 15-2.
  - Elapsed turn-on time is up to 12 user days.
  - One administrator can only handle 300-500 users.

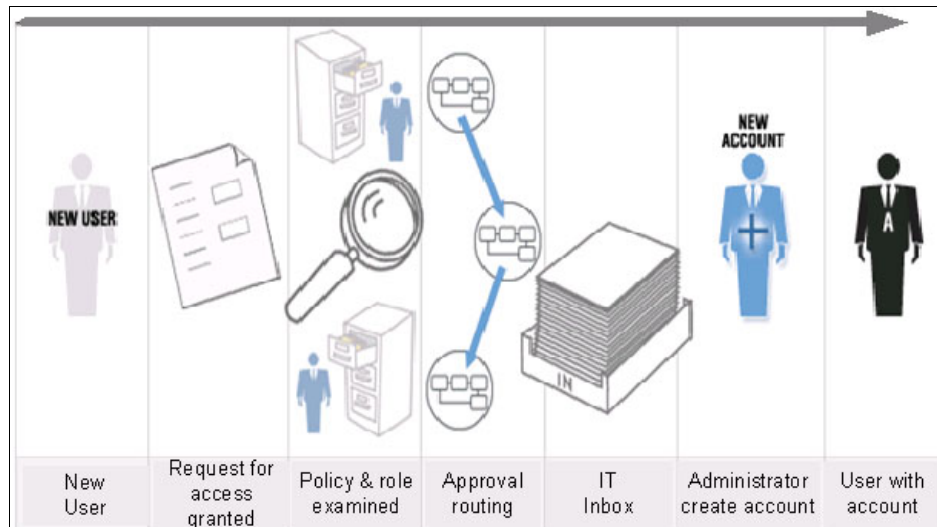


Figure 15-2 Problem with provisioning new user

- **Problem 2:** Users have multiple user accounts and passwords. This problem is illustrated in Figure 15-3.
  - Each help desk call costs \$20-\$45.
  - Up to 60% of help desk calls are password resets.

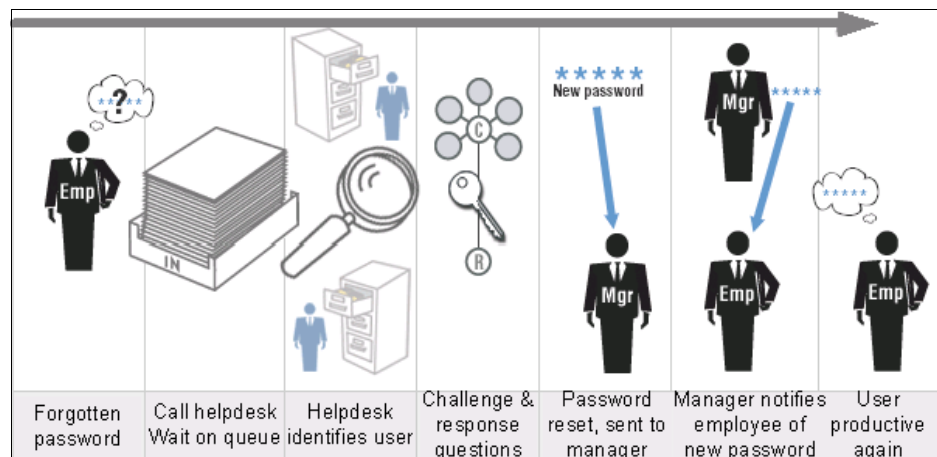


Figure 15-3 Problem when managing user accounts

- **Problem 3:** ITSO-Electronics.com uses slow and inconsistent processes to revoke user accounts and de-provision user access rights, as shown in Figure 15-4.
  - Account turn-off performance is poor. At any one time, 30 to 60% of accounts are invalid.
  - One administrator only handles 300-500 users.

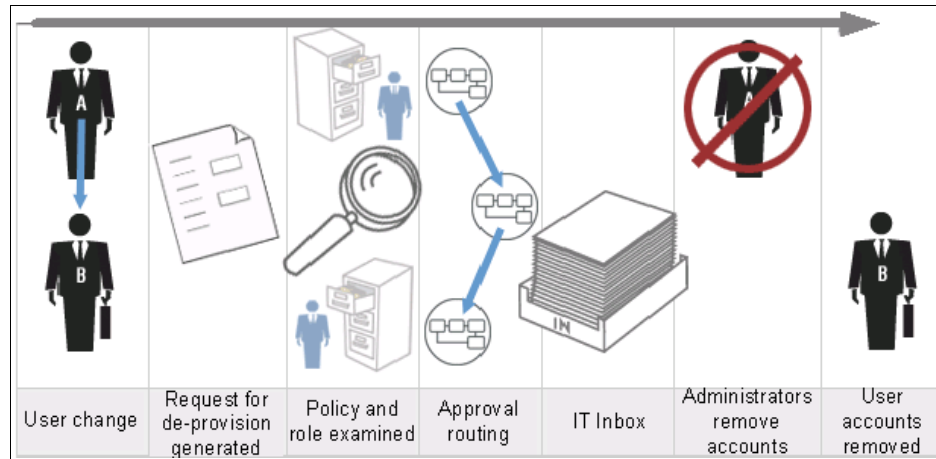


Figure 15-4 Problem of de-provisioning users

- **Problem 4:** Application and data security is custom-written into each business initiative. This problem is illustrated in Figure 15-5.
  - Custom security costs \$40-\$80K to develop per application.
  - Users have unique security identities per application.

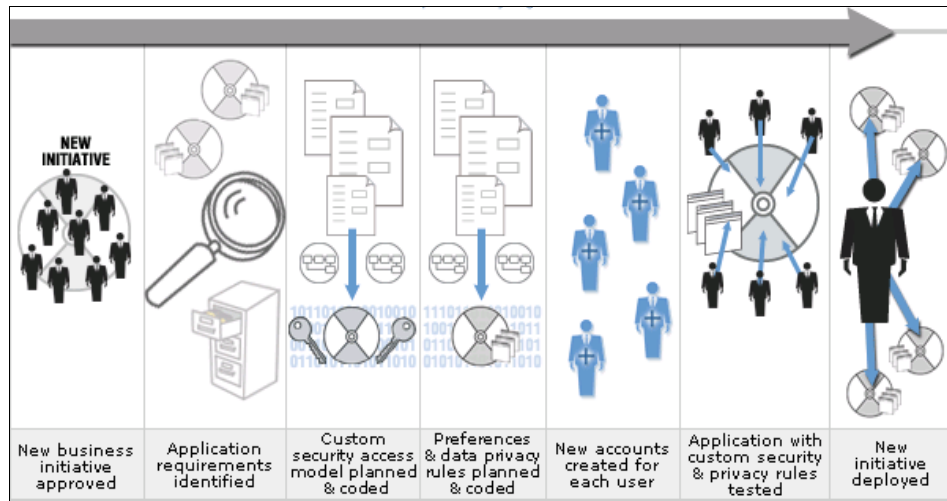


Figure 15-5 Problem with new initiatives

## Inconsistent security architecture

Figure 15-6 summarizes the existing security architecture deployed by ITSO-Electronics.com, with multiple Web server host systems deployed in the Internet DMZ




Figure 15-6 Current ITSO-Electronics.com architecture

As seen in Figure 15-6, each application has its own security control and architecture. The most pressing issues, summarized from four main points of view, are as follows:

***From a business point of view***

- ▶ High cost and risk of entry
- ▶ Low utilization
- ▶ Vertical integration offers little flexibility
- ▶ Multiple, competing security infrastructures
- ▶ Long-term maintenance staffing is required

***From a security and technical point of view***

- ▶ Limitations in scalability
- ▶ The security model is too operating system centric
- ▶ Key components are exposed within the DMZ
- ▶ It is difficult to apply a uniform security model
- ▶ It is difficult to keep up with evolving standards
- ▶ Authentication is not flexible for requirements

***From the user administration point of view***

- ▶ Multiple administrators with multiple access control tools
- ▶ Tools do not work together
- ▶ User and access control information is everywhere

***From the user point of view***

- ▶ Too many passwords to remember

### 15.4.5 Solution approach

We have described the company and technical background of ITSO-Electronics.com. We also talked about the pressing issues, challenges, and difficulties that ITSO-Electronics.com is currently facing.

IBM's integrated Identity Management solution can help ITSO-Electronics.com get users, systems, and applications online, productive, and secure fast, reducing costs and maximizing return on investment (ROI). This suite of solutions automates and simplifies the management of user identities, access rights, and privacy policies across the e-business infrastructure.

The integrated IBM Tivoli software solutions can leverage ITSO-Electronics.com's current systems, eliminating redundant components and improving business interactions. Processes are automated, information access decision-making is unified, and regulatory requirements are satisfied, resulting in lower costs, a more secure posture, passed audits, enhanced IT service levels, and a better user experience.

Figure 15-7 shows various capabilities of an identity management solution.

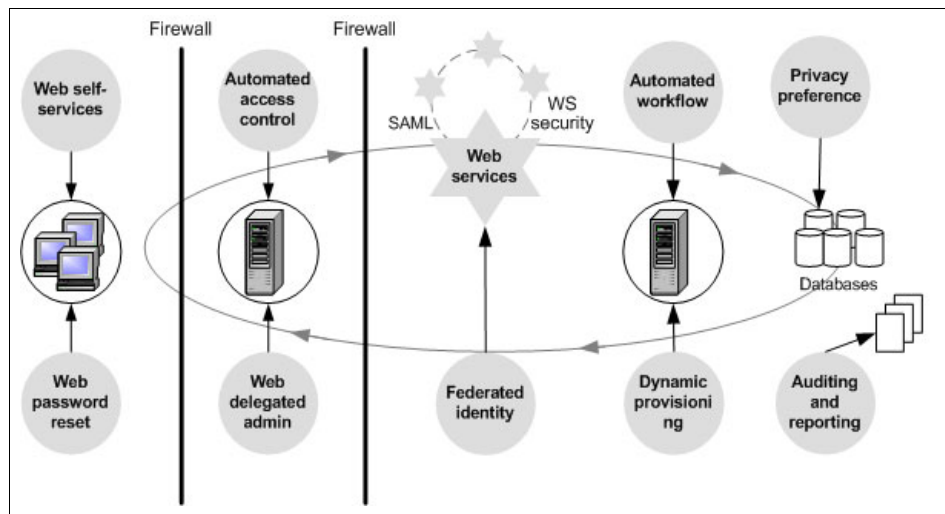


Figure 15-7 Capabilities of an identity management solution

To help ITSO-Electronics.com to effectively manage a broad range of business constituencies, IBM Tivoli Integrated Identity Management addresses four key areas of identity management (Figure 15-8).

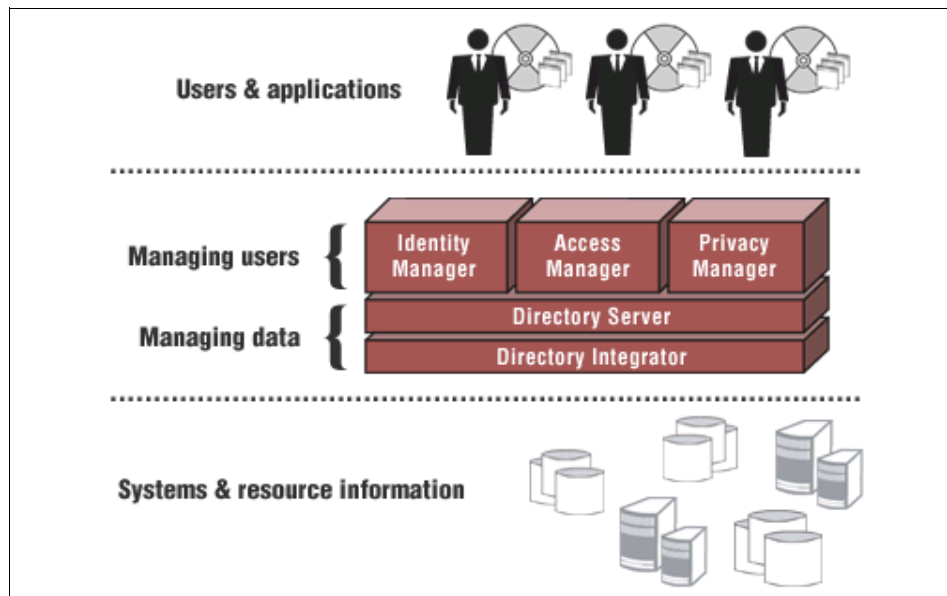


Figure 15-8 IBM's Integrated Identity Management solution

### ***User provisioning***

IBM Tivoli Identity Manager centrally coordinates enrollment, provisioning, and user self-care tasks. This streamlines and automates workflow, increasing accuracy and efficiencies, and reducing administrative and operational costs.

### ***Access control***

The suite of IBM Tivoli Access Manager solutions (Tivoli Access Manager for Business Integration, Tivoli Access Manager for e-business, and Tivoli Access Manager for Operating Systems) provides identity control through access control, single sign-on, and auditing, reducing administrative costs and simplifying the administration of data protection and access control policies.

### ***Privacy control***

IBM Tivoli Privacy Manager for e-business protects consumer trust and brand integrity by implementing privacy policies which control sensitive information. Centrally developing and deploying privacy policy rules enables adherence across the IT infrastructure, creating efficiencies and reducing administration costs.



### ***Identity synchronization***

IBM Directory Solutions (IBM Tivoli Directory Server and IBM Tivoli Directory Integrator) provide an authoritative, identity data infrastructure to serve as a base for advanced security and Web applications.

### **Implementation of IBM Identity Management solution**

The following sections summarize the implementation of this solution.

#### ***Initial architecture***

The decision of how many components of a security solution to apply to a specific enterprise varies from one company to another. Although an IBM Tivoli Identity Management solution can address all four key areas in Identity Management, we will be using three of the four key components to address the current issues for ITSO-Electronics.com.

#### ***Identity synchronization***

IBM Tivoli Directory Server is a powerful Lightweight Directory Access Protocol (LDAP) engine that serves as the identity data infrastructure for deploying comprehensive identity management applications and advanced software architectures. Built on industry-leading DB2 for optimum reliability and scalability, IBM Tivoli Directory Server provides the high performing and high availability identity data repository that maximizes the value of key directory-enabled applications such as portals and identity management.

IBM Tivoli Directory Server provides the efficiency of open architecture that enables this solution to fit into the widest variety of heterogeneous environments without disruption, regardless of current operating system or directory. Standards and Web Services support enable these solutions to integrate seamlessly with a wide variety of repositories and technologies and provides integration with both existing and new Web Services within an enterprise.

#### ***User provisioning***

IBM Tivoli Identity Manager is a software solution that solves ITSO-Electronics.com's user provisioning issues with policy-based identity management that works across legacy and e-business environments. It allows ITSO-Electronics.com to manage user roles and responsibilities centrally, enabling ITSO-Electronics.com to realize lower costs, higher levels of security, and a rapid return on investment.

Tivoli Identity Manager features centralized Web administration, providing a consistent security implementation across the entire organization while simplifying management through a single interface. Its embedded provisioning engine and universal integration tools makes the job of managing users more efficient by quickly provisioning them with appropriate resources and reducing

the administrative workload. Its auditing and reporting mechanisms will quickly produce reports for internal audits, ensuring regulatory compliance. Importantly, this solution will help the enterprise operate more efficiently and reduce costs through:

- ▶ Automating processes for account creation, change requests, and approvals
- ▶ Eliminating inaccuracies in user permissions
- ▶ Coordinating existing systems to improve accuracy
- ▶ Providing role-based delegated administration
- ▶ Providing a self-service interface to ease the burden of daily administration on help desk and IT staff

### ***Access control***

The IBM Tivoli Access Manager solutions (Access Manager for Operating Systems, Access Manager for Business Integration, and Access Manager for e-business) provides consistent identity access control from a single administration console, which consolidates management, as well as simplifies administration of access control policies. These solutions provide centralized access enforcement across all Web applications, Microsoft Windows desktop applications, UNIX operating systems, and WebSphere MQ. With single sign-on to these applications, sign on and password confusion is eliminated, making access enforcement more efficient and unburdening overloaded help desk resources.

The solution also provides flexible single sign-on from browsers, PCs, and wireless devices to any Web server, using flexible authentication such as passwords, certificates, or biometrics. Reusable security components for all custom applications mean that policy is defined once and then applied consistently across all applications, saving development time and money.

Together with WebSphere, IBM Tivoli Access Manager for e-business can provide an integrated solution for e-business security needs, combining core security technologies with common security policies. This provides a business with:

- ▶ Single sign-on across multiple applications
- ▶ Lowered help desk costs and improved customer satisfaction
- ▶ Central repository for authorization policy, helping to reduce administration costs
- ▶ Delegated administration of user management and authorization policy, allowing clear delineation of administrative responsibilities
- ▶ Support for multiple authentication mechanisms, including certificates, without affecting the target application

The initial architectural diagram is shown in Figure 15-9.

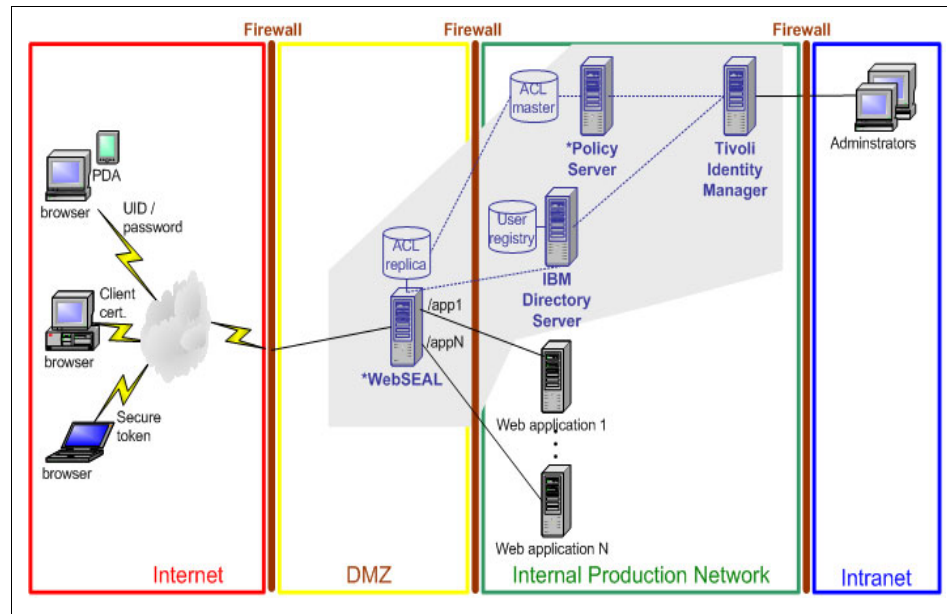


Figure 15-9 Initial IBM Identity Management architecture

WebSEAL and Policy Server are two key components of Tivoli Access Manager for e-business.

## Connecting the pieces

Now that the key components are all in place, this section describes how the components interact with one another.

The Internet-facing WebSEAL will be listening on ports 80 and 443 (SSL). We will also modify the configuration of the Web application server in the internal production network slightly to listen on alternate ports; in our case, we use ports 81 and 1443. This permits ITSO-Electronics.com to close the common ports 80 and 443 on the firewall between the DMZ and production networks.

Access to common LDAP ports (389/636) is also disallowed from the Internet, because WebSEAL is the only entity that needs to communicate from the DMZ to the user registry.

There is also the question of whether the junctions between the Internet-facing WebSEAL and the Web application servers require the use of SSL. In this case, because the Web application servers are in a controlled zone, it is not strictly necessary. If the Web application servers were in the open corporate intranet,

SSL should probably be used. The choice to use SSL may be made based upon the specific risk associated with the content involved. The answer is similar with respect to communication with the user registry.

### **Instant benefits**

As highlighted in Figure 15-8, the components of the IBM Identity Management solution (including IBM Directory Server, Tivoli Identity Manager, WebSEAL and Policy Server) form the “Security On Demand Infrastructure” for ITSO-Electronics.com. This initial architecture provides the following instant benefits:

- ▶ The security model is independent of the operating system.
- ▶ A limited component (only WebSEAL as the reverse proxy) has exposure within the DMZ.
- ▶ It is architecturally consistent and we have a uniform security model.
- ▶ It is not dependent on internal resources to support core security component code.
- ▶ As standards evolve, the security infrastructure can be readily upgraded.
- ▶ Since various Web applications are shielded by the DMZ WebSEAL, users can enjoy instant single sign-on capability between various Web applications through WebSEAL.
- ▶ Administrators can manage both user management and access control with one unified Web-based administration tool for various Web applications.

### **Providing internal employee access**

The initial Security On Demand Infrastructure is focused on providing access for outside Internet users. There is also a need for ITSO-Electronics.com to allow its employees to access the same Web applications in the internal production network from the corporate intranet. In other words, Internet application access is currently only being provided for client applications and content.

ITSO-Electronics.com could route employee browser traffic to internal applications through the same WebSEAL that resides in the Internet DMZ. However, this is not a recommended approach, partly for security reasons, and partly for manageability and performance reasons.

This could be easily accomplished by putting in another WebSEAL server that is dedicated solely to internal access (shown in Figure 15-10). This allows ITSO-Electronics.com to create a different set of junctions for the internal and external WebSEAL servers, which permits better segregation of content between the two access classes.

**Tip:** There is an interesting issue here that we will touch on briefly, but not dwell on. That is, there may be scenarios where it makes sense to have different user namespaces for employees and clients. This can easily be accomplished by creating a second Access Manager secure domain. However, in this scenario, such requirements do not exist. In this architecture, we will keep it simple and use a single Access Manager user registry covering both employee and client users in a common user ID namespace.

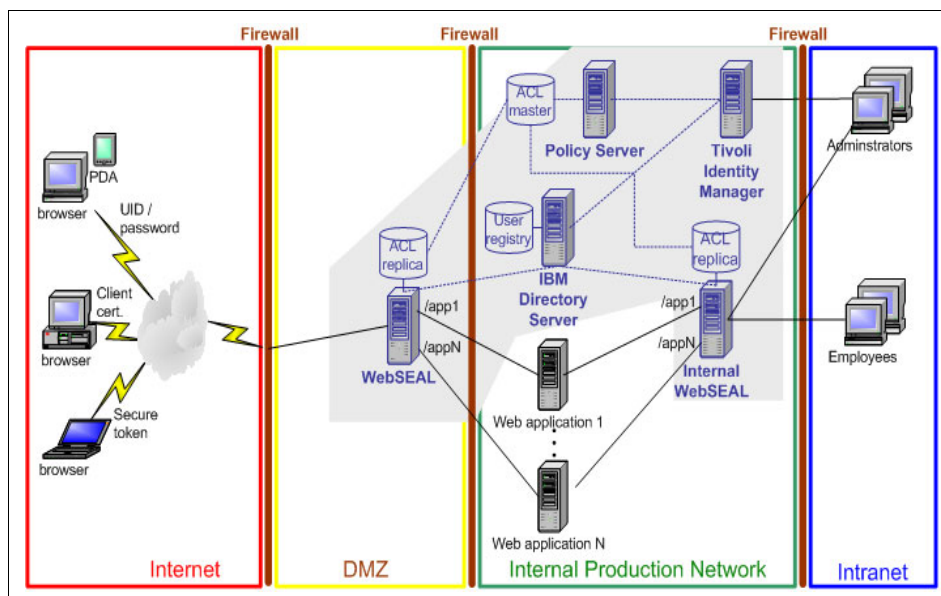


Figure 15-10 Security architecture with internal WebSEAL

### Scalability, performance, high availability, and resilience

Because of the anticipated growth in the number of Internet users (an annual growth of 50 percent is expected over the next five years), a robust, highly scalable, and high availability system/architecture is required. The following discussion shows how easily these on demand requirements can be achieved with the Security On Demand Infrastructure in ITSO-Electronics.com.

Scalability is the ability to respond to increasing numbers of users who access resources in the secure domain. High availability and resilience ensures the system provides 24 x 7 services and that user requests can be handled evenly by load balancing to the available services for instant response.

Tivoli Access Manager for e-business uses the following techniques to provide scalability:

- ▶ Replication of services:
  - Authentication services:
    - The IBM Directory Server (LDAP) distributed architecture supports scalable directory services with server replication capabilities. Server replication improves the availability of a directory service. Tivoli Access Manager treats the LDAP server replication as a master-subordinate relationship. The combination of a master server and multiple replicated servers helps ensure that directory data is always available when needed. If any server fails, the directory service continues to be available from another replicated server.
    - Tivoli Access Manager supports this replication capability by connecting to the LDAP master server when it starts up. If the LDAP master server is down for any reason, the Tivoli Access Manager server must be able to connect to an available LDAP replica server for any read operations.
    - Many operations, especially those from regular users, are read operations. These include operations such as user authentication and signon to back-end junctioned Web servers. After proper configuration, Tivoli Access Manager performs failover to a replica server when it cannot connect to the master server.
  - Authorization services:

Authorization service components such as WebSEAL can be replicated to increase availability in a heavy-demand environment. Multiple mirrored authorization services are used to load balance user requests to ensure high performance, high availability, and resilience. There is no master and slave concept in Tivoli Access Manager's authorization services. Each authorization service is independent from the others. Should any part of one of the authorization services fails, another available mirrored authorization service will take up the work; hence, no failure as a whole.
  - Security policies:

The master authorization policy database, containing policy rules and credential information, can be configured to automatically replicate. The master authorization policy database is responsible for all the update operations of access control while the replicated policy database resides with the authorization services to provide rapid access control query requests.

- ▶ Front-end (WebSEAL) replicated servers:  
Mirrored resources for high availability and resilience. Load balancing client requests are handled by the mirrored WebSEALs.
- ▶ Back-end replicated servers:
  - Back-end servers can be WebSEAL or third-party application servers.
  - Mirrored resources (unified object space) for high availability.
  - Additional content and resources.
  - Load balancing of incoming requests through junctions.
- ▶ Optimized performance by allowing the off-loading of authentication and authorization services to separate servers
- ▶ Scaled deployment of services without increasing management overhead

## **Scaling the Security On Demand Infrastructure**

The following sections describe the various components to be added into the initial architecture.

### ***Additional WebSEAL servers***

An additional DMZ WebSEAL server is added into the environment to mirror the existing DMZ WebSEAL to accept and process more Internet browser traffic. Should the demand for Internet access rise rapidly, additional (mirrored) WebSEAL servers can be easily joined into the secure domain to ensure the Security On Demand Infrastructure is highly scalable.

With additional WebSEAL servers to handle the Internet traffic, performance should have sufficient improvement.

### ***Directory server replicas***

In order to handle the growth of the user base, additional directory server replicas can be added into the secure domain to improve performance, and also provide high availability and resilience to the security infrastructure.

### ***Network dispatcher***

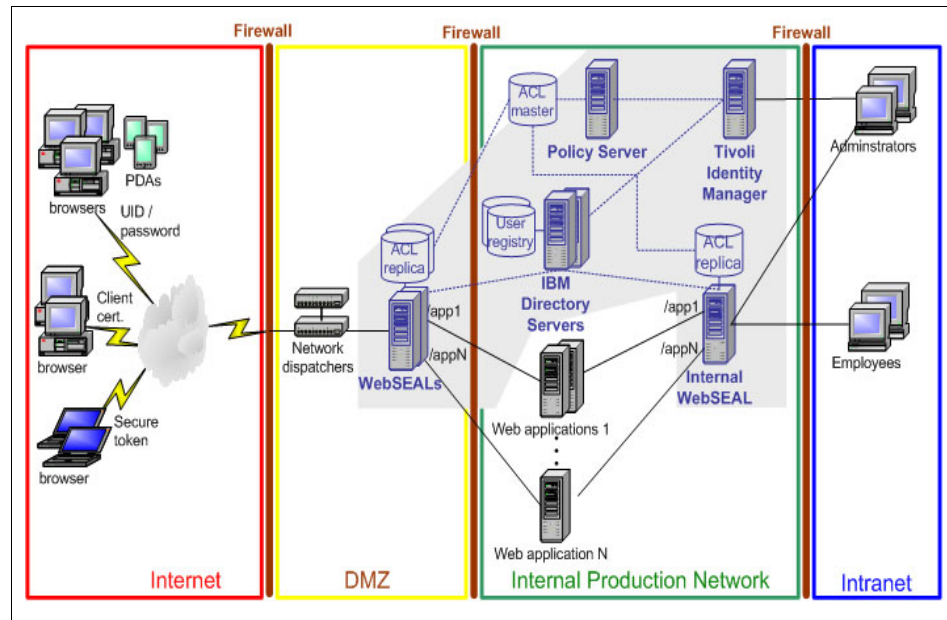
One or more network dispatchers, provided by the WebSphere Edge Server component of the WebSphere Application Server Network Deployment capability, are placed in front of the two DMZ WebSEAL servers. The network dispatcher acts as the front door to load balance the workload evenly between the two DMZ WebSEALs.

It also has the ability to detect if there is any failure of one of the DMZ WebSEALs. Should this happen, all traffic will be automatically routed to the next available WebSEAL to ensure system availability.

### ***Additional back-end Web application servers***

If more power is required in the back-end Web application server, ITSO-Electronics.com can simply set up additional Web application servers to mirror the existing ones and add them to the same junction point to the DMZ and Internet WebSEALs.

The final security infrastructure architectural diagram is shown in Figure 15-11.



*Figure 15-11 Security architecture with high availability, scalability, and resilience*

#### **Points to note in Figure 15-11:**

1. Additional internal WebSEALs can be added into the infrastructure as needed. There is no restriction on how many mirrored WebSEAL or authorization services can be added into the secure domain.
2. There should only be one “Policy Server” within a secure domain. The Policy Server is responsible to interact with Web Portal Manager and Tivoli Access Manager administration command prompts for user registry and policy update. Authentication and authorization services are not affected if the Policy Server is unavailable.



## **New internal-only highly sensitive Web application**

The Human Resource (HR) department of ITSO-Electronics.com is going to put their employee information online so that all employees in ITSO-Electronics.com can update their personal information by themselves. The HR department also wants to set up an online collaboration teamroom for employees to express their ideas and suggestions for the company. As a result, a new HR Web application called “Employee Online” is developed by the IT department of ITSO-Electronics.com.

However, the sensitivity of employee information is always a big concern and issue of the HR department. The “Employee Online” Web application requires high security. The data access level also varies between different parties. For example, an HR manager may have access to all employee information, while an HR clerk may only have access to general employee information. Employees themselves may have access to their own employee information while they are not allowed to access any other employee’s information.

Using the traditional way ITSO-Electronics.com fulfilled these complicated and dynamic information access control requirements, the IT department needs a lot of resources to study, design, implement, test, and deploy such an application. User access control management and system maintenance may also require extensive resources and staff. Using the Security On Demand Infrastructure already in place can reduce the demand this puts on the IT and HR departments. The time and effort required to deploy such an internal Web application is minimal.

Following is a summary of what needs to be done to satisfy the HR requirement:

- ▶ A new WebSEAL instance is set up in the internal production network. This new WebSEAL instance can be resident on the same machine with the existing internal WebSEAL server, hence, no additional hardware is required.
- ▶ To enforce strict access control to the system, client certificate authentication is set up in the new WebSEAL instance, and only HTTPS traffic is accepted. Each employee will be issued a client certificate installed in their workstation to identify themselves to the “Employee Online” system. A password is set by the employee to protect the client certificate.
- ▶ The “Employee Online” Web application is placed in the internal production network. An HTTP junction is set up between the new WebSEAL instance and the “Employee Online” Web server because both new WebSEAL instances and the “Employee Online” Web application are located within the trusted internal production network. There is definitely no access from the outside Internet world. Employees can only access the system from their own workstations from their internal office.

- ▶ The application development team makes use of the Tivoli Access Manager authorization application programming interface to allow “Employee Online” to query the authorization service to make fine-grained authorization decisions.

The authorization API is the interface between the resource manager (requesting the authorization check) and the authorization service itself. The authorization API allows the policy-enforcing application “Employee Online” to ask for an authorization decision, but shields the application from the complexities of the actual decision-making process.

The authorization API provides a standard programming model for coding authorization requests and decisions. The authorization API lets the application team make standardized calls to the centrally managed authorization service from “Employee Online.”

There are two modes of authorization API that can be used: Remote cache mode and Local cache mode.

- **Remote cache mode:** In this mode, the API is initialized to call the remote authorization server to perform authorization decisions on behalf of the application. The authorization server maintains its own cache of the replica authorization policy database. This mode is best suited for handling authorization requests from application clients.
  - **Local cache mode:** In this mode, the API is initialized to download and maintain a local replica of the authorization database for the application. Local cache mode provides better performance because the application performs all authorization decisions locally instead of across a network. However, the overhead of database replication and the security implications of using this mode make it best suited for use by trusted application servers.
- ▶ The HR manager and Risk Management (RM) manager will work together to define the roles and access rules for the “Employee Online” system. Once all the roles and rules are defined, they will pass the initial access control list (ACL) to the service delivery group to deploy. The ACL is frequently reviewed by the HR manager, RM manager, and auditors. Should any changes need to be made, the ACL will be passed to the system administrator to execute. This enforces separation of duties to ensure the system is highly secure and carefully monitored.

The new security infrastructure is depicted in Figure 15-12.

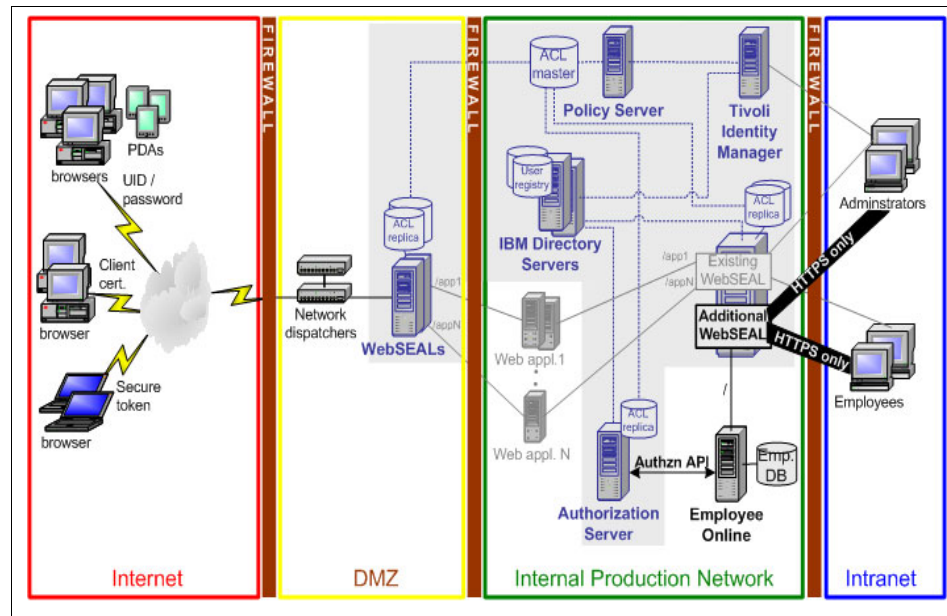


Figure 15-12 New internal and highly sensitive Web application

**Note:** There is no linkage between the DMZ WebSEALs and the “Employee Online” Web application. This ensures that no outside Internet access to enforce high security and protection for the employee information.

## User provisioning

Security requirements within a business are becoming more and more dynamic in order to react to market changes in a timely manner. Every day, every minute, every second, there may be new users, existing users changing roles, group membership updates, and deletion of obsolete users in a dynamically provisioned environment. These changes are required to be reflected within the infrastructure and synchronized so that the defined security policies can take effect immediately.

At the same time, granting or denying access rights to users and enforcing new policies to enable runtime decisions in real time should be exercised at once.

We have talked about requirements on the application side. Now, we focus on how the Security On Demand Infrastructure handles user and policy provisioning.

These tasks only require well-defined processes between the related departments and the system administrators. Neither hardware nor software, nor any system configuration changes are required.

### ***New user or employee***

When a new user or employee joins ITSO-Electronics.com, a system administrator can create the new user or employee record in the user registry via Tivoli Identity Manager and add the new user or employee to the related groups/roles. The new user will instantly receive the required privileges to access the appropriate systems.

For example, Linda is a new Human Resource (HR) employee. She is responsible for calculating employee payrolls. She needs to access the Employee Online system to update her own personal information. She also requires access to other employees' salary-related information.

There are two related groups set up in ITSO-Electronics.com's Security On Demand Infrastructure system, namely, stocks4u-employees and stocks4u-payroll. The policy rules associated with these two groups are as follows:

- ▶ Any users within the group stocks4u-employees have the rights to access the Employee Online system and their own personal information.
- ▶ Users within the group stocks4u-payroll group have the rights to access other employees' salary-related information.

By adding Linda to the ITSO-Electronics.com directory service and putting her into the stocks4u-employees and stocks4u-payroll groups, Linda can now authenticate herself to the Security On Demand Infrastructure through the internal WebSEAL. If Linda is requesting to access Employee Online to update her personal information, the internal WebSEAL will allow her to do so based on her group membership in stocks4u-employee and the associated policy defined in the policy database. The same applies if Linda wants to access another employee's salary information; WebSEAL will grant her access since she is a member of the group stocks4u-payroll.

### ***Change of roles for user or employee***

When a user or employee changes their roles within ITSO-Electronics.com, system administrators can remove the user or employee from one group or role and add them to another related group or role. Please note that one user or employee can be added to more than one group/role as they may have more than one role assigned within ITSO-Electronics.com.

Let's take Linda's case as an example. After six months, Linda's manager decided to assign Linda to an additional role as the HR consultant to provide advice and support to other employees. One of the additional tasks that Linda is required to do is to record all of the employee questions to the Web application called "Employee FAQ". This application is associated with the company's database. ITSO-Electronics's employees are free to access this Employee FAQ application while only the editor of this application has the right to update the database.

In order to allow Linda to update the records in the Employee FAQ database, administrators can add Linda to the pre-defined group "stocks4u-employee-faq-editor". The associated policy with this group is that members within this group can edit the database. As a result, Linda can update and record on any employee's queries into this database.

### ***Change of company policy***

ITSO-Electronics.com used to allow only its customers to access its research Web site "Funding Online." This research Web site provides online stock and fund inquiries for ITSO-Electronics.com customers. It also has a forum for ITSO-Electronics.com customers to exchange thoughts and ideas.

With the rapid growth of competition between ITSO-Electronics.com and other brokerage companies, ITSO-Electronics.com wants to open up this research Web site to the public so that everyone can access the site, while only registered customers can use the forum facility. By opening up this research Web site to the public, ITSO-Electronics.com can more easily communicate to not just its customers, but the public as a whole, especially possible future clients.

To open up the public access to "Funding Online" while limiting access to its discussion forum to registered customers only, ITSO-Electronics.com needs to perform the following steps:

- ▶ Define and apply a new access control list (ACL) in the master policy database within the Security On Demand Infrastructure that any unauthenticated users are allowed to access the public contents of the "Funding Online" Web site.
- ▶ Refine and apply the ACL that only authenticated users (someone registered within ITSO-Electronics.com's user registry but not required to be with any group membership) can access the discussion forum.

As a result, everybody can access ITSO-Electronics.com's "Funding Online" research Web site, but only registered ITSO-Electronics.com customers can use the discussion forum facilities.

### ***User self-registration and self-care***

It would be resource intensive to manage user accounts, especially to add new users to the ITSO-Electronics.com directory service, once the “Funding Online” research Web site is opened up to the public. ITSO-Electronics.com foresees that there could be more than 500 new users to be added into the corporate user registry on a daily basis.

Tivoli Identity Manager includes a sample application that allows end-users to perform self-registration. Self-registration is the process by which a user can enter required data to become a registered user, without the involvement of an administrator. Users enter specific identification information (either company specific or user specific) with a user ID and password. The identification information provided by the user is then validated and the user is created in ITSO-Electronics.com user registry.

Tivoli Identity Manager also supports self-care. Users are able to go to the Web delegate administration page and manage their passwords. After logging into the delegate administration page, users can go to the Change My Password task to manage their password.

### ***Denying access of user or employee to a particular Web application***

As time goes by, Linda’s workload is getting heavier and heavier. She decides to tell her manager that she cannot handle both payroll and employee queries tasks. After further discussion with her manager, she is reassigned to do the employee consultation only.

So, by removing Linda from the group “stocks4u-payroll”, Linda no longer has access rights to other employee’s salary information. WebSEAL will deny any request from Linda to access other employee’s salary information. This is an example of how denying access to a user or employee can be simply achieved.

### ***Revoking existing user or employee accounts***

There will be cases in ITSO-Electronics.com that existing users or employees may leave the company. When such a departure occurs, the only task that the System Administrator must perform is to simply remove the user or employee ID from the user registry.

Say after one year, Linda accepts a new job in Company ABC, and resigns from ITSO-Electronics.com. Her identity within ITSO-Electronics.com’s directory server is permanently removed so that Linda has no access to any system within ITSO-Electronics.com.

## Summary of IBM Identity Management solution

As a result of implementing the IBM Identity Management solution described in this chapter, ITSO-Electronics.com can develop Web Services applications that are secured beyond the firewall. This means they can facilitate secure transactions with partners in their supply chain, regardless of the Web Services or security technologies used by other partners. See Figure 15-13.

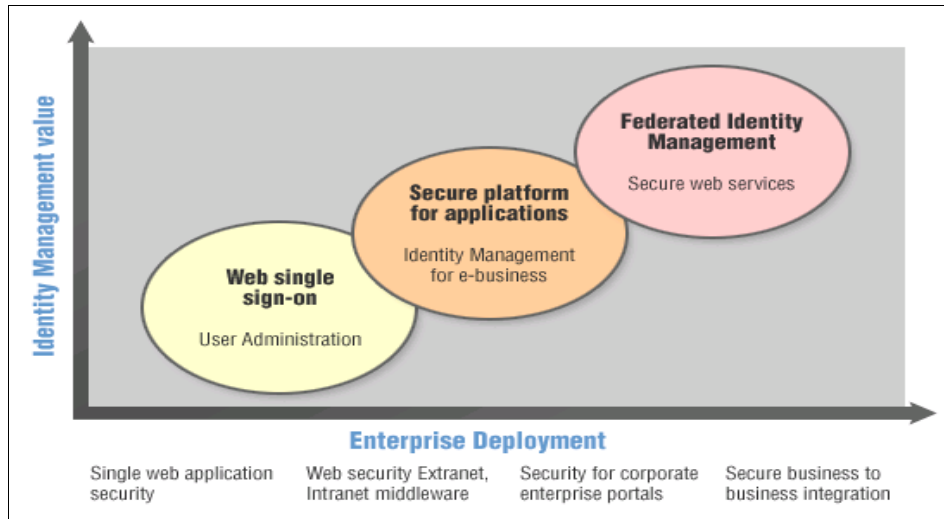


Figure 15-13 Shifting value of identity management

### **Authentication infrastructure**

The authentication infrastructure eliminates the need for application developers to develop code to authenticate users. The company does not need to invest in developing any authentication capabilities within its new applications.

### **Authorization infrastructure**

The authorization infrastructure provides the ability to control access to Web applications and content, which may be hosted through multiple Web servers, at the URL-level. It also provides the ability to make fine-grained authorization decisions within applications via the Tivoli Access Manager authorization API.

### **High scalability, high performance, high availability, and resilience**

The infrastructure has the ability to respond to increasing numbers of users who access resources. It ensures that the system provides 24 x 7 services, while user requests can be handled evenly by load balancing to the available servers for instant response.

### ***Single sign-on capability***

Tivoli Access Manager for e-business WebSEAL provides the single sign-on capability for all ITSO-Electronics.com Web-based applications. A user should only need to log in one time to obtain access to all authorized applications and content which may reside on various servers. The growth in e-commerce, business-to-business, and Web Services initiatives has driven companies to provide secure access to Web applications for business partners, customers, and employees.

### ***Cross-platform security solution***

Previous experience with the in-house security application clarified the need to maintain operating system independence for Web-based application security. With the Security On Demand Infrastructure, security can be enforced across different platforms. The application servers can be on UNIX, Linux, Windows, AIX, O/S 400, and z/OS.

### ***Zero client footprint for accessing applications***

The solution supports browser-based access to applications for both employees and customers. From their desks, internal users can access both Internet-hosted applications and internal applications.

### ***User provisioning***

The solution provides e-provisioning tools to mediate among human resources systems, directory services, and network resources to automatically set up new user accounts in the applications that employees need to do their jobs. These tools also automate de-provisioning by removing user accounts from the system when users or employees no longer need the resources. The tools allow ITSO-Electronics.com to accomplish both user and policy tasks far more quickly, securely, and cheaply than they can with manual processes.

### ***Centralized and delegated user and access control management***

There is a centralized tool for administrators to manage users, groups, permissions, and policies. This Web-based tool has the extensibility beyond delegated user management to also deliver delegated security administration.

## **15.5 Product positioning**

The IBM products that address security as described in this approach are the IBM Tivoli Identity Manager, the IBM Tivoli Access Manager, IBM Tivoli Directory Server, the IBM Tivoli Directory Integrator, and the IBM Tivoli Storage Manager.

In the following sections, we discuss some aspects of secure access and how the referenced IBM products provide solutions today.



## 15.5.1 Identity management

Figure 15-14 represents an identity management blueprint.

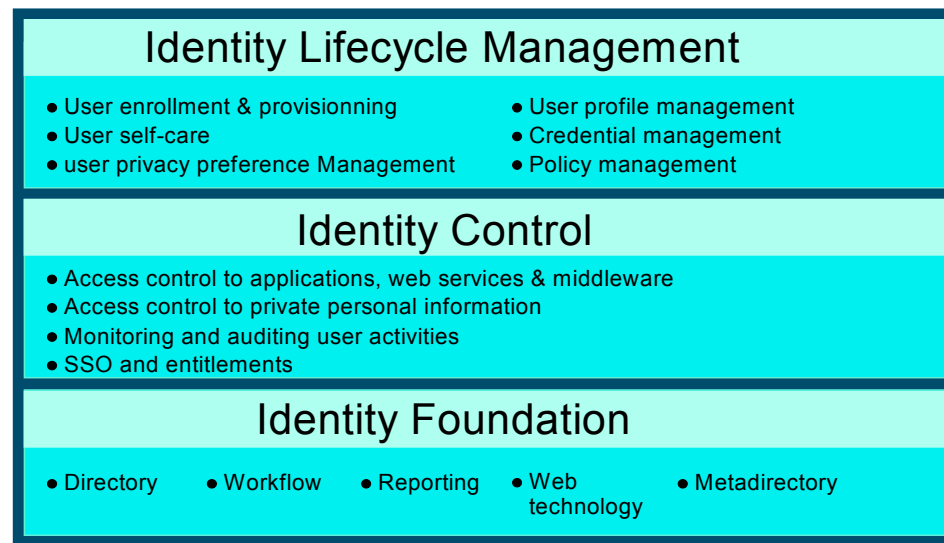


Figure 15-14 Identity management blueprint

### Identity lifecycle management

The first layer, identity lifecycle management, is addressed by the IBM Tivoli Identity Manager product. Tivoli Identity Manager has several capabilities that make it a flexible and extensible solution to manage user information, including:

- ▶ **Simple, ubiquitous interfaces:** Tivoli Identity Manager provides ease-of-use through Web interfaces for end-user self-care and administration commands. These firewall-friendly interfaces are accessible from a Web browser.
- ▶ **Out-of-the-box integration:** Tivoli Identity Manager can support existing environments with ready access to a broad set of endpoints and many common reports. Other integration points include interfaces with HR systems, help desk and IBM Tivoli Access Manager. These integration points can help relieve developers from having to support these interfaces, and thereby speed deployment.
- ▶ **Models for expandability and customization:** An identity management solution needs to adapt to its IT environment. The scalable data capabilities of Tivoli Identity Manager enable administrative flexibility by using technologies such as delegated administration, role-based access control structures, and toolkits for custom endpoint support and expansion of the GUI, workflow, and reporting.

IBM Tivoli Identity Manager also provides:

- ▶ Self-care
- ▶ Password change/reset
- ▶ Registration and enrollment Web interfaces
- ▶ Delegated administration
- ▶ Role-based access-control structures
- ▶ Automated attribute generation and validation
- ▶ Metadirectory connectivity
- ▶ Toolkits for GUI
- ▶ Endpoint support
- ▶ Workflow
- ▶ Reporting

## **Identity control**

The identity control layer is made up of IBM Tivoli Access Manager and IBM Tivoli Privacy Manager. IBM Tivoli Access Manager for Operating Systems securely locks down business-critical applications, files, and operating platforms to help prevent unauthorized access. This security capability blocks both insiders and outsiders from unauthorized access to and use of valuable customer, employee, and business partner data. Additionally, Tivoli Access Manager for Operating Systems audits application and platform activity. This capability provides the assurance customers, employees, and partners expect.

Tivoli Access Manager addresses the requirements of an On Demand Operating Environment by providing the following capabilities:

- ▶ Defense against the top security threat that enterprises face: misbehavior by internal users and employees
- ▶ Delivery of mainframe-class security and auditing in a lightweight, easy-to-use product
- ▶ A combination of full-fledged intrusion prevention — host-based firewall, application and platform protection, user tracking and controls — with robust auditing and compliance checking
- ▶ Persistent Universal Auditing to document compliance with government regulations, corporate policy, and other security mandates

## **Identity foundation**

The bottom layer in the identity management framework is addressed by two products: IBM Tivoli Directory Server and IBM Tivoli Directory Integrator.

### ***IBM Tivoli Directory Server***

IBM Tivoli Directory Server provides a powerful Lightweight Directory Access Protocol (LDAP) identity infrastructure that is the foundation for deploying comprehensive identity management applications and advanced software architectures such as Web Services, including the following features and capabilities:

- ▶ LDAP V3 support ensures compatibility with industry standard LDAP-based applications.
- ▶ Reliable IBM DB2 Universal Database V8.1 engine provides scalability to tens of millions of entries, as well as groups of hundreds of thousands of members.
- ▶ Robust replication capability for both master/subordinate replication, gateway, cascaded, and peer-to-peer replication with up to dozens of master servers.
- ▶ Easier management and usability with a Web administration GUI and features such as dynamic and nested groups, along with sorted and paged search results, tight integration with IBM operating systems, WebSphere middleware, and Tivoli identity management and security products.
- ▶ Broad platform support: AIX, Solaris TM, Microsoft Windows 2000, and HP-UX TM, as well as Linux distributions for Intel, and IBM @server iSeries, pSeries and zSeries platforms.

### ***IBM Tivoli Directory Integrator***

The IBM Tivoli Directory Integrator synchronizes identity data residing in directories; databases; collaborative systems; applications used for human resources (HR), customer relationship management (CRM), and Enterprise Resource Planning (ERP); and other corporate applications.

With some built-in connectors, an open-architecture, Java development environment to extend or modify these connectors, and tools to apply logic to data as it is processed, IBM Tivoli Directory Integrator can help by:

- ▶ Synchronizing and exchanging information between applications or directory sources
- ▶ Managing data across a variety of repositories, providing the consistent directory infrastructure needed for a wide variety of applications including security and provisioning
- ▶ Creating the authoritative data spaces needed to expose only trustworthy data to advanced software applications such as Web Services

## Enterprise Identity Mapping

Today's network environments are made up of a complex group of systems and applications, resulting in the need to manage multiple user registries. Dealing with multiple user registries quickly grows into a large administrative problem that affects users, administrators, and application developers. Consequently, many companies are struggling to securely manage authentication and authorization for systems and applications. Enterprise Identity Mapping (EIM) is an IBM @server infrastructure technology that allows administrators and application developers to address this problem more easily and inexpensively than previously possible

This is an architecture for describing the relationships between individuals or entities (such as file servers and print servers) in the enterprise and the many identities that represent them within an enterprise. This architecture provides a common set of APIs that can be used across platforms to develop applications that look up the relationships between user identities and a single EIM identifier that represents a user in the enterprise.

For example, given a person's user identity in one user registry, one can determine which user identity in another user registry represents that same person. If the user has authenticated with one user identity and one can map that user identity to the appropriate identity in another user registry, the user does not need to provide credentials for authentication again. Therefore, EIM provides a generalized identity mapping function for the enterprise. Applications now have the flexibility of using one user registry for authentication while using an entirely different user registry for authorization.

No code changes are required to existing user registries. The administrator does not need to have mappings for all identities in a user registry. EIM allows one-to-many mappings and many-to-one mappings. EIM is an open architecture that administrators can use to represent identity mapping relationships for any registry of any type. It does not require copying existing data to a new repository and trying to keep both copies synchronized. The only new data that EIM introduces is the relationship information. Administrators manage this data in an LDAP directory, which provides the flexibility of managing the data in one place and having replicas wherever the information is used.

EIM provides the mechanics for cross-platform user identity management.

One can write applications that use EIM APIs to perform look-up operations for user identities within an enterprise. Each of these APIs is supported across all the IBM @server platforms.

The main components of the EIM architecture are:

- ▶ **EIM domain controller:** An LDAP server configured to manage at least one EIM domain.
- ▶ **EIM domains:** A directory within an LDAP server contains EIM data for an enterprise. An EIM domain is the collection of all the EIM identifiers, EIM associations, and user registries that are defined in that domain.
- ▶ **EIM identifiers:** Represent a person or entity in an enterprise.
- ▶ **EIM registry definitions:** Represent an actual user registry that exists on a system within the enterprise.
- ▶ **EIM associations:** Describe a relationship between an EIM identifier that represents a specific person and a single user identity in a user registry that also represents that person.

The EIM domain controller can be hosted by:

- ▶ AIX on pSeries with IBM Directory V5.1
- ▶ Linux with IBM Directory V5.1
- ▶ OS400 on iSeries with OS400 V5R2 Directory Services
- ▶ Windows2000 on xSeries with IBM Directory V5.1 client
- ▶ z/OS on zSeries with z/OS V1R4 LDAP

EIM Administration tools are available on:

- ▶ OS400 on iSeries with iSeries Navigator V5R2
- ▶ z/OS on zSeries with V1R4 LDAP SPE OW57137

## 15.5.2 Privacy Control Management

IBM provides multilevel security support in z/OS Version1 Release 5. Multilevel security addresses government requirements for highly secure data which can be shared between agencies. While multilevel security began as a government requirement, it is now apparent that this new technology has applications in the general business sectors as well, as security controls become more critical in the on demand, virtual environments. A multilevel security system has two primary goals:

- ▶ First, the controls are intended to prevent unauthorized individuals from accessing information at a higher classification than their authorization.
- ▶ Second, the controls are intended to prevent individuals from declassifying information.

Multilevel security function will allow customers more stringent access control to resources than is provided by user permissions.

Designed together with DB2 Version 8, a solution is available now for multilevel security on the zSeries mainframe. New security features in DB2 V8 and z/OS 1.5 enable customers to have a single repository of data which can be accessed by different agencies, by people with different need-to-know authority. This highly secure access is managed at the row/column level in DB2 to provide the granularity that is required. Each row of a DB2 database can now contain a DB2 security label. This label is checked against the security label of the requester. Only rows containing matching security labels can be accessed by the requester.

The security of a row or column is maintained at a Database Administration or Security Administration level. An application developer does not need to worry about adding security-sensitive code to a Query or Insert. For example, one doesn't need to have a Query on `xyx` where `security_level="TopSecret."` The developer would only need to code, Query on `xyz`. The security "test" is implicitly executed by the combination of DB2 V8 and z/OS security server with RACF®.

Earlier, labeling was provided and MVS™ received B1 security certification already in 1990. In z/OS V1.5 (available since September 2004), multilevel security extended the labeled security protection of z/OS to include TCP/IP and UNIX System Services. This enables z/OS to meet the stringent requirements for Multilevel Security.

The robust mainframe capabilities of z/OS and zSeries servers for high availability and high scaling (through Parallel Sysplex and GDPS, for example), make this an excellent candidate for handling the multilevel labeling requirements of customers.

## 15.6 Linkages

Table 15-1 shows various IBM products and their features that can be linked together to form a solution in an On Demand Operating Environment.

Table 15-1 IBM product usage and linkages

Area	Product name	Usage
User provisioning	IBM Tivoli Identity Manager (ITIM)	Manages user roles and responsibilities.
	*IBM Directory Server (IDS)	User registry for ITIM.
	IBM WebSphere MQ	User provisioning workflow within ITIM.
	IBM WebSphere Application Server (WAS)	Hosts the ITIM application.
Access control	IBM Tivoli Access Manager (ITAM)	Provides cross-grained and fine-grained access control for the back end Web resources.
	WebSEAL	One of the major components of ITAM, it acts as the Web reverse proxy to provide authentication and authorization for the back-end Web application server.
	Authorization Server	One of the major components of ITAM, it performs authorization decisions on behalf of the application.
	*IBM Directory Server (IDS)	User registry for ITAM.
Identity synchronization	IBM Directory Integrator (IDI)	Synchronizes identity data residing in directories, databases, and collaborative systems.
Application server	IBM WebSphere Application Server	Back-end Web application server.

**Notes:**

IBM Directory Server (IDS) can be shared between IBM Tivoli Access Manager (ITAM) and IBM Tivoli Identity Manager (ITIM).

Do not confuse IDS, IBM Directory Server, with Intrusion Detection Services.

Figure 15-15 shows a high-level overview of how the products are related to each other. As can be seen here, there are overlaps between the products in each area. We will further explain each component in the following sections.

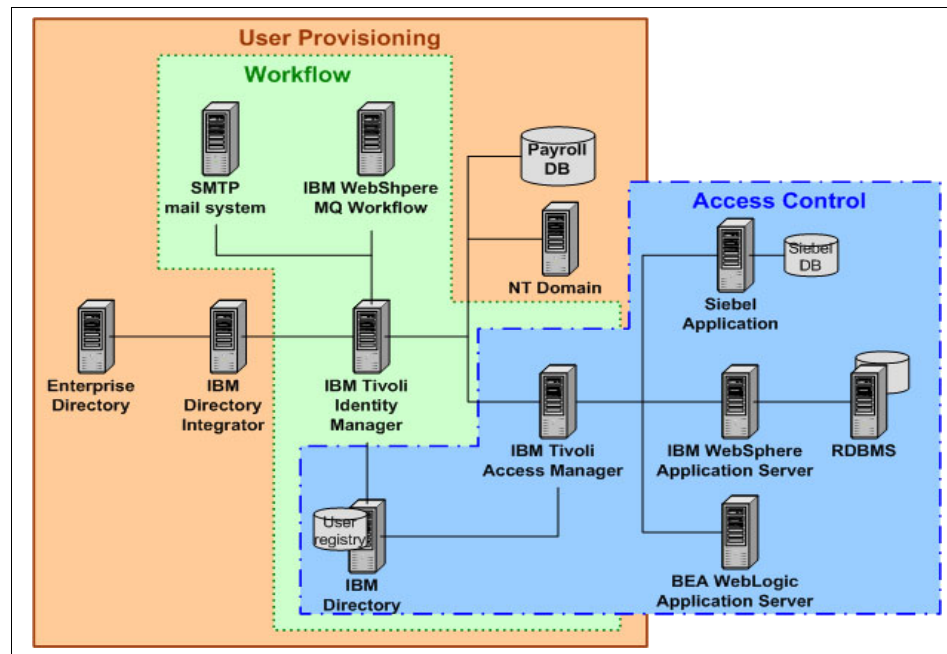


Figure 15-15 Product linkage overview

### **Workflow**

Entitlement workflows define the business logic that is tied specifically to provisioning actions. Since Entitlement workflows are part of the provisioning process, they must have a final process state of APPROVED or REJECTED. If the entitlement workflow has an APPROVED process state, the provisioning operation will be executed. If the entitlement workflow has a REJECTED process state, the provisioning process stops and the provisioning request has a status of failed.

IBM Tivoli Identity Manager leverages the workflow capabilities of IBM WebSphere MQ Workflow.

Figure 15-16 demonstrates how a request for user access permission is handled by IBM Tivoli Access Manager.



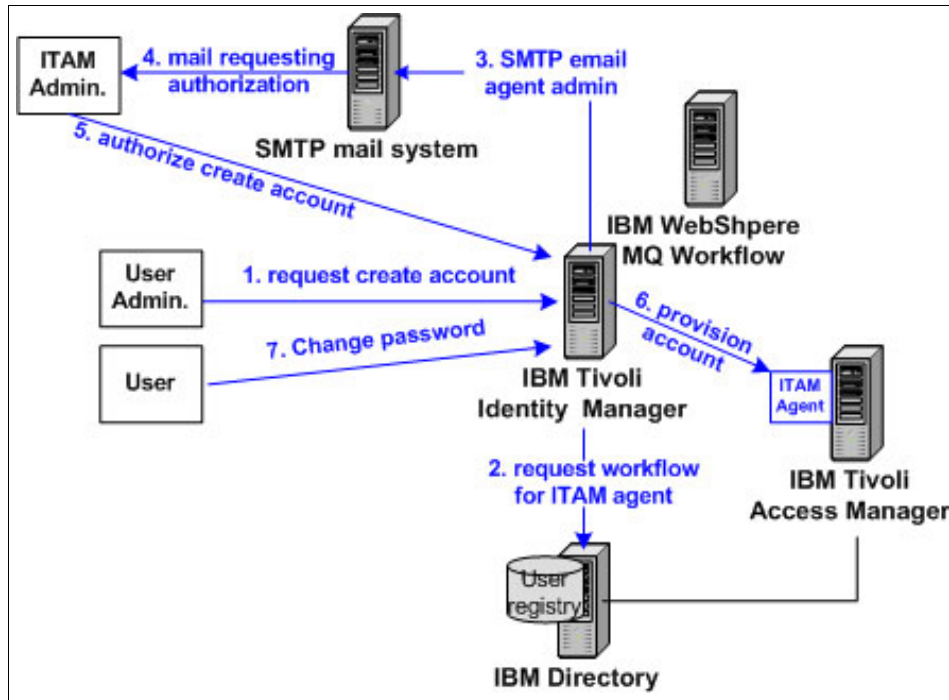


Figure 15-16 User provisioning workflow

The workflow pictured in Figure 15-16 is described as follows:

1. A user administrator requests a user account creation in the IBM Tivoli Access Manager domain.
2. IBM Tivoli Identity Manager recognizes that a workflow is configured for the provisioning of accounts to the IBM Tivoli Access Manager service. It then queries the IBM Directory for the workflow details. IBM Tivoli Identity Manager uses IBM WebSphere MQ Workflow to enforce the workflow process.
3. IBM Tivoli Identity Manager sends a request to the SMTP mail system to notify the IBM Tivoli Access Manager administrator to authorize the new user account creation request.
4. The IBM Tivoli Access Manager administrator receives an e-mail that contains a secure HTTP link to the related IBM Tivoli Identity Manager approval page.
5. The IBM Tivoli Access Manager administrator can accept or reject the user account creation request by pressing the “APPROVE” or “REJECT” button.

6. Upon receiving the “APPROVED” response from the IBM Tivoli Access Manager administrator, IBM Tivoli Identity Manager then provisions the account by sending a Darpa Agent Markup Language (DAML) create command to IBM Tivoli Access Manager agent.
7. The user is now able to manage their IBM Tivoli Access Manager account with IBM Tivoli Identity Manager.

### ***User provisioning***

In case a new entry is manually added to the Enterprise directory with the “changelog” feature enabled, a changelog event is sent to the IBM Directory Integrator. IBM Directory Integrator then performs an LDAP lookup with the attributes in the changelog entry. If the object is found, IBM Directory Integrator sends a JNDI feed to IBM Tivoli Identity Manager.

IBM Tivoli Identity Manager uses the information gathered from the IBM Directory Integrator to create a new user within the IBM Tivoli Identity Manager domain. It then automatically provisions the new accounts with its authorized agents. As shown in Figure 15-17, a new IBM Tivoli Access Manager user, an NT user, and a Payroll user are created within its domain.

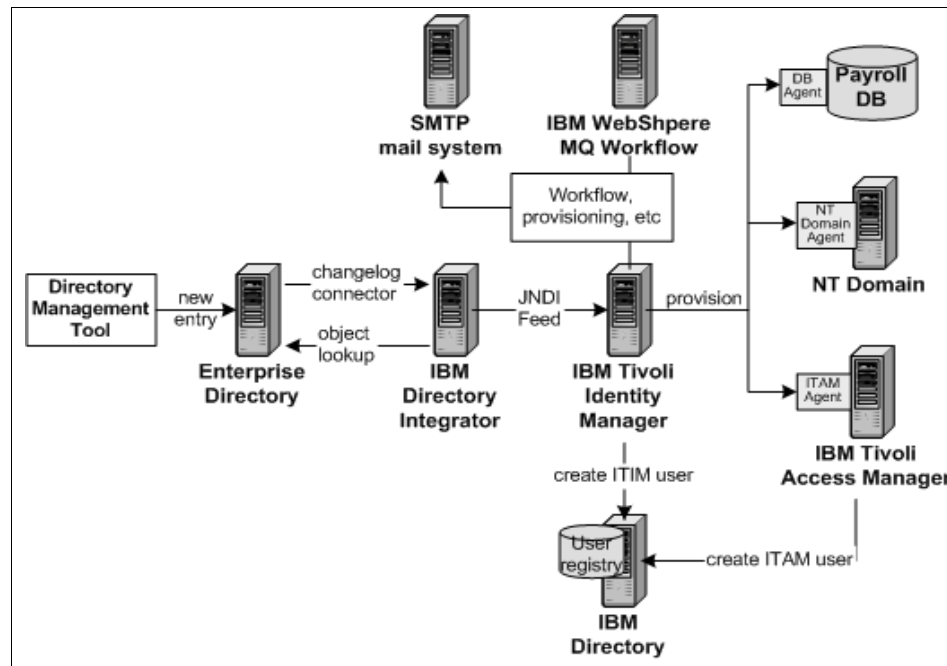


Figure 15-17 User provisioning

### ***Access control and single sign-on capability***

Figure 15-18 shows how authentication, access control, and single sign-on can be achieved within the IBM Tivoli Access Manager secure domain. WebSEAL acts as the one and only entry to all of the back-end protected applications. All external user requests are required to be passed via WebSEAL first. Each user is required to have one user identity and password to authenticate with WebSEAL. Once the user is logged into IBM Tivoli Access Manager via WebSEAL, they can access the various back-end applications according to their granted privileges.

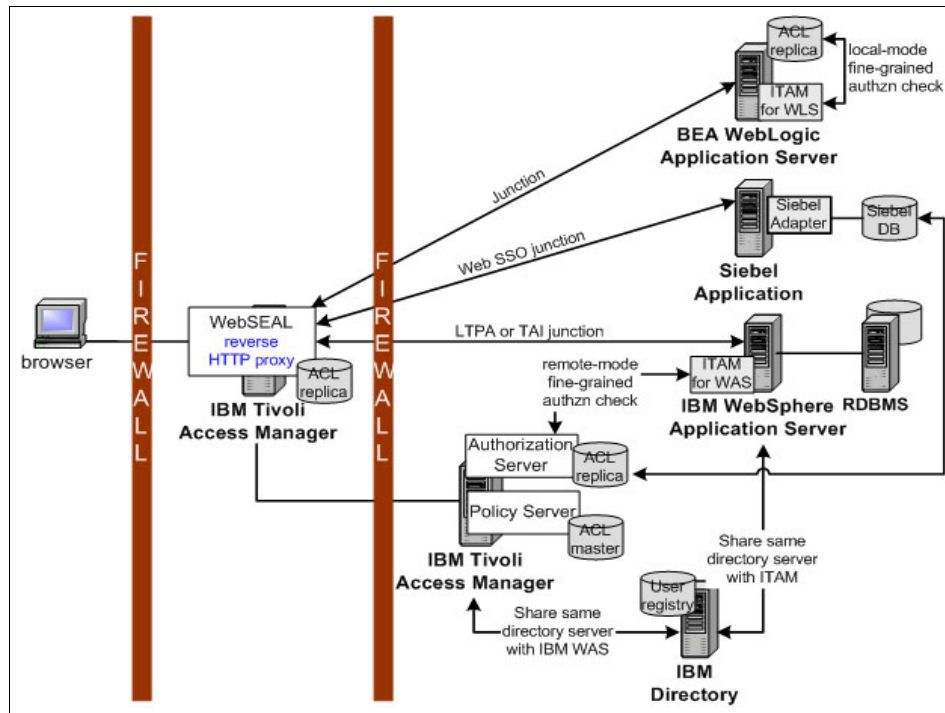


Figure 15-18 Access control and single sign-on to various back end applications

### ***WebSEAL and WebSphere Application Server***

There are two ways to link between WebSEAL and IBM WebSphere Web Application Server: Lightweight Third-Party Authentication (LPTA) or Trust Association Interceptor (TAI). TAI is the preferred interface.

**LPTA** LPTA is an authentication mechanism that enables multiple or heterogeneous servers to share authentication information during a session; thus, the user does not have to re-authenticate when accessing another server.

**TAI** TAI is a mechanism that lets a server forward user credentials to the WebSphere Application Server.

When a WebSphere application with J2EE security is run, and the user tries to access a protected resource, WebSphere authenticates the user by consulting the user registry. The user registry is shared with Tivoli Access Manager.

When the user requests access to a protected method or resource, the WebSphere container uses information from the J2EE application deployment descriptor to determine the required role membership.

The WebSphere container uses the integrated Tivoli Access Manager module to request an authorization decision (“granted” or “denied”) from the Tivoli Access Manager authorization server.

The WebSphere container also passes additional context information, when present, to the authorization server. The optional context information can include cell name, host name, and server name. If the Tivoli Access Manager policy database has policies specified for any of the context information, the authorization server can use this information when making the authorization decision.

The authorization server consults the Tivoli Access Manager user definitions in the shared user registry. The authorization server then consults the permissions that have been defined for the specified user within the Tivoli Access Manager protected object namespace. The Tivoli Access Manager authorization server returns the access decision to the WebSphere container.

WebSphere Application Server either grants or denies access to the protected method or resource.

### **Tivoli Access Manager and Siebel eBusiness applications**

The IBM Tivoli Access Manager adapter for Siebel eBusiness Applications maps the Siebel security model into the Access Manager security model. IBM Tivoli Access Manager’s WebSEAL component provides Web single sign-on and a single point of control for securing Siebel eBusiness Applications. User profiles, groups, and role definitions are stored in IBM Directory (LDAP). WebSEAL implements a secure HTTP/S reverse proxy that is configured in the demilitarized zone. Siebel servers are located in the secure network behind the inner firewall.

At runtime, HTTP or HTTPS sessions are intercepted by WebSEAL, which checks to see if the request is for a secure Siebel resource that is protected by IBM Tivoli Access Manager’s security policy. After successful authentication and authorization, WebSEAL builds the Web SSO credentials and maps the user privileges into the Siebel application roles and credentials. These are then passed to the Siebel application, which presents the right set of application capabilities to the user based on the user’s roles.

## **WebSEAL and BEA WebLogic application server**

An external user requests access to a protected resource. The request is received by WebSEAL before entering the secure network of the enterprise. WebSEAL authenticates the user in the Tivoli Access Manager secure domain. WebSEAL applies its own authorization decision based on the requested URL and the Tivoli Access Manager access policy.

WebSEAL can apply considerations such as account validity, time-of-day, and authentication mechanism. After the user's request has been authorized, WebSEAL forwards the request to the WebLogic server. The request includes the external username and a special password within the basic authentication header. The special password belongs to the "configured user," and allows the Tivoli Access Manager Custom Realm to confirm WebSEAL as the origin of the request.

The WebLogic server transparently passes the authenticated user identity and password to the Tivoli Access Manager Custom Realm. The Tivoli Access Manager Custom Realm uses Tivoli Access Manager authentication services to verify that the password provided by WebSEAL is correct for the "configured user" mentioned previously. That is, this password provides the basis of trust that the request's origin is WebSEAL.

When a request for a J2EE resource is received by WebLogic Server, it checks the relevant deployment descriptor information to determine if access to the resource is restricted to certain roles. If the request requires the user to assume a role, the WebLogic Server queries the Tivoli Access Manager Custom Realm to determine whether the requesting user is a member of any of the groups that are mapped to the role. The Tivoli Access Manager Custom Realm consults the Tivoli Access Manager authorization server to determine if the current user is a member of the group.

If the user is a member of a group that is mapped to a permitted role, access is granted. Otherwise, access is denied.

## 15.7 Glimpse of the future

Some of the directions that we see security taking in the future will enhance the On Demand Operating Environment and bring even more benefits to enterprises that are prepared to take advantage of them.

- ▶ The first one would involve hardware and operating systems with the adoption of The Trusted Computer Platform Alliance (TCPA), now transitioning to Trusted Computing Group (TCG). This emerging standard has the goal of certifying a platform (both hardware and software) as being trusted as it boots up. Once trusted, performing transactions with this platform becomes easier.
- ▶ A second advance would empower the user administration with real federated identity management, where users declared on one computer would have single sign-on capability to every system they need to access.
- ▶ Another involves inter-application security; this is where the Web Services security standards come into play. When an application tries to access a Web Service, it will have the possibility to request a security token. This token could be a Security Assertion Markup Language (SAML) authorization token issued by a trusted organization, or this could also be a more traditional certificate, based on keys. IBM has already announced that it is reworking its software portfolio to be more modular and take this new protocol into account.

Tivoli Access Manager v4.1 and v5.1 support SAML and Web Services.

The IBM's Federated Identity Management (FIM) solution extends identity management for both the identity provider and service provider infrastructure. IBM Tivoli federated identity management solution builds on the current Tivoli identity and security offerings. In a federated environment, a user can log on through his identity provider in order to conduct transactions or easily access resources in external domains.

Partners in a federated identity management environment depend on each other to authenticate their respective users and vouch for their access to services. Federated identity standards, such as those being produced by the Liberty Alliance or the Web services security specifications, form an encapsulation layer over local identity and security environments of different domains. This encapsulation layer provides the ingredients for interoperability between disparate security systems inside and across domains, thus enabling federation.

For more information, please refer to the redbook *Federated Identity Management with IBM Tivoli Security Solutions*, SG24-6394.

## 15.8 Summary

In this chapter we have described a business issue centered around secure access to systems, applications, and data. We have described some of the products available today that could apply to businesses wanting to address similar issues. We have also described a scenario where some of these products are used together to provide a solution.

Taking a focused approach to meeting today's business needs by adopting products and solutions that fit into the On Demand Operating Environment framework allows customers to evolve their environment in a step-wise fashion. They can adopt individual components of the overall framework as needed to address current business requirements, knowing that the On Demand Operating Environment will also be able to adapt to changing needs and future requirements, allowing them to be responsive to the needs of their business and their customers.







## How to provision system resources according to business demands

This chapter focuses on the need for efficient and automated provisioning and workload management to support an On Demand Business. To roll out new applications and adjust the resources allocated to specific business applications, systems need to be provisioned in a fast and reliable way. A key component of the On Demand Operating Environment is virtualization. One aspect of virtualization is the transparency to the user and the application of the underlying physical resources that are being used and accessed.

By enabling an environment where resources can be dynamically re-provisioned to meet the immediate needs of the business, enterprises can become more responsive and resilient. Dynamic re-provisioning allows unused or under-used resources to be applied to applications that may be experiencing a spike in demand. This not only helps meet immediate business requirements but saves costs associated with over-provisioning hardware for anticipated, but infrequent spikes for specific applications.

This chapter identifies business needs as they relate to provisioning system resources in an On Demand Operating Environment and identifies the products available today from IBM that can be used in a solution to address these

business needs. The products are mapped to the On Demand Operating Environment framework to show where they fit. A practical scenario is used to illustrate challenges that IT organizations face regarding provisioning or re-allocating of resources.

## 16.1 Introduction

Provisioning in an On Demand Operating Environment is the capability to automatically deploy and dynamically optimize operational resources in response to business objectives in a heterogeneous environment. There are many types of resources that may be provisioned: servers, distributed systems, storage, identities, applications, networks, and more.

When platform provisioning technologies are combined with other provisioning technologies and enhanced by orchestration, the On Demand Operating Environment will have the ability to make the most informed decisions about provisioning hardware, software, applications, and so on to optimize the IT infrastructure.

Businesses need to be able to address the following requirements:

- ▶ Quickly deploy new N-tier applications and all of the resources that support them to stay responsive and competitive
- ▶ Support a heterogeneous environment, with multiple hardware platforms, operating systems, and software and middleware components.
- ▶ Reduce administrative costs associated with deployments through the use of automation.
- ▶ Increase system utilization by quickly and reliably re-provisioning systems to meet immediate business requirements including service level objectives.
- ▶ Make the underlying hardware and OS platforms as transparent to the applications, business processes, and users as possible.

## 16.2 General strategy

The components needed to implement this approach have to respect the general attributes of an On Demand Operating Environment. Specifically, they must be:

- ▶ **Self-managing:** To reduce both administrative overhead associated with provisioning systems and human error, the provisioning process should be as automated as possible, and self-managing. This implies the use of autonomic technologies to sense when systems need to be re-provisioned, and taking advantage of automated workflows to carry out the re-provisioning steps.

- ▶ **Scalable:** Solutions need to be scalable to be able to address the provisioning of systems across large server farms and cluster environments. As companies merge with one another, an On Demand Operating Environment would allow for the integration of new systems and applications, and also enable the re-provisioning of systems that may no longer be required due to redundancy.

Provisioning solutions that take advantage of the existing provisioning technologies provided with many hardware and software platforms are inherently scalable to some extent, by having the capability to offload some of the work to the individual platform.

- ▶ **Resilient:** A provisioning solution can be a major enabler of a resilient environment by allowing for quick provisioning of systems to recover from catastrophic failures. But the provisioning system itself must also be resilient, able to recover from errors during the provisioning process.
- ▶ **Economical:** Any solution today needs to be economical, and a provisioning solution is no different. A centralized solution that addresses multiple platforms and takes advantage of the unique facilities provided by some of those platforms takes advantage of economies of scale. In addition, it reduces the administrative overhead.
- ▶ **Open standards-based:** The only way that a centralized solution can provide consistent provisioning to a large number of platforms, operating systems, applications, and data stores is by building on and using open standards.

## 16.3 Solution components

This approach focuses on both the virtualization and the automation components of infrastructure management. It covers the basic hardware resources and how servers, storage, and networks must become virtualized. To support applications, computing resources need to be adapted to support varying loads. Optimally, this should occur transparently to the application so that the application does not have to be stopped and restarted in order to take advantage of new resources allocated to it.

Regarding servers, an On Demand Operating Environment must include systems that have the capability to adapt their processing power to the actual need, for instance, increasing or decreasing the number of processors, speed, or memory allocated based on the current workload. A workload manager component would be used to dynamically adjust the resources to meet business requirements.

The number of systems allocated to a function may also need to be changed to meet requirements. In this case, we need an efficient way to install the correct operating system as well as the correct set of middleware and the required applications on the newly added node or nodes.

Grid computing can be applied to transparently use otherwise under-utilized systems. Job scheduling and load balancing of complex tasks are critical capabilities.

From a storage standpoint, the requirements are similar. The goal here is to be able to expand storage capacity and have the operating systems able to dynamically recognize it and translate this into more space for file systems and ultimately the applications that require it.

The provisioning of storage within the Tivoli Storage Resource Manager is based upon Filesystem Extension Policy created by the local system administrators. The policy that can be set is based on a number of factors which determine the scope of the automation and for which resources this automation applies.

Provisioning of networks requires both hardware and middleware capabilities since the operating environment should be able to identify new network adapters and add them to an existing IP environment. Having virtual IP addresses may be the most efficient way to achieve this transparently for the application.

All these points describe how the IT resources can react to a change in the environment. But the operating environment also needs ways to monitor these changes, and provide a clear view of the resources used, in terms of CPU, memory, storage, or network utilization. Automation, combined with virtualization, enables the optimization of those resources so several applications, previously supported in their own silos, can exchange resources dynamically and with little or no manual effort on the part of system administrators.

In order to take this action, the operating environment must be able to map enterprise policy to actual IT resources. It must have a complete inventory of the existing IT resources and what they are used for at any point in time.

## 16.4 Scenario

In this section, we continue our discussion of the same fictitious company, “ITSO-Electronics.com,” that was described in the previous chapter.

## 16.4.1 Current environment

As described in “Data centers” on page 173, the San Diego data center acts as the entry point for all Internet application access and key internal application access. There are two major applications that reside in this data center:

- ▶ **Employee Online:** Employee Online is a highly sensitive Internal-only Web application used by the Human Resource (HR) department for all employee-related matters such as payroll calculations, employee benefits calculations, year-end bonus calculations, vacation calculations, and so on.

This Web application is also being used by all employees within ITSO-Electronics.com for employee information updates and inquiries, vacation tracking, and various employee benefit claims.

- ▶ **Stocks Online:** Stocks Online is an Internet Web application that provides open access to the public for general company and stock information inquiries.

Stocks Online is also an online trading site that has been designed to make it easier for ITSO-Electronics.com customers to manage their investments. The online trading service is available only to existing clients who deal directly with ITSO-Electronics.com to buy, sell, and switch stocks and funds, as well as to obtain information on their current holdings.

Both Web applications are protected by the IBM Integrated Identity Management solution described in 15.4.5, “Solution approach” on page 181.

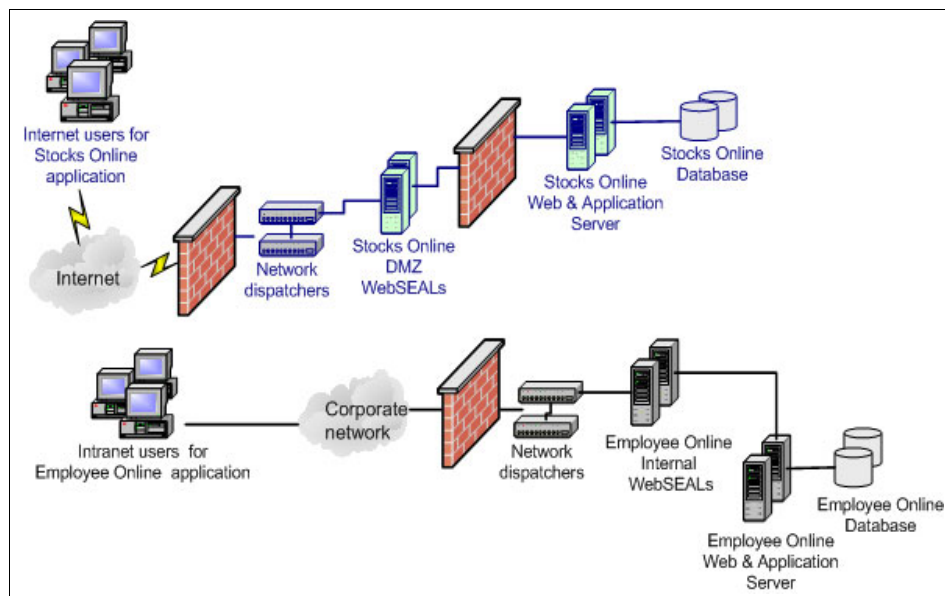


Figure 16-1 Current Stocks Online and Employee Online applications

## **Stocks Online application**

Stocks Online is a business-critical application which has resource requirements with predictable fluctuations, specifically, high demand on weekdays during the Stock Market operating hours, at the end of each month, and at the end of each financial year.

System stability is highly desired as ITSO-Electronics.com customers cannot bear any delay or downtime of the system that may potentially cause financial loss to them. As a result, high availability and resilience are the most important requirements during the peak hours. In general, Stocks Online requirements are driven by customer demand.

Briefly, the Stocks Online application components are as follows:

### ***DMZ WebSEAL***

DMZ WebSEAL is the security component that resides in the DMZ as the secure Web reverse proxy to protect Stocks Online (refer to “Implementation of IBM Identity Management solution” on page 183 for more details about WebSEAL). It consists of multiple servers with the following software:

- ▶ Linux Red Hat 7.3
- ▶ IBM Tivoli Access Manager for e-business WebSEAL 4.1

Additional servers are provisioned for the WebSEAL component due to the business-critical nature of the Stocks Online application.

### ***Stocks Online Web and application server***

The Web and application server function has several servers running with the following software:

- ▶ AIX 5.1
- ▶ Apache HTTP Server 1.3
- ▶ IBM WebSphere Application Server 5.0.2
- ▶ Stocks Online IBM WebSphere Application Server component

### ***Stocks Online database***

The Stocks Online database server is run with the following software:

- ▶ z/OS V1R4.0
- ▶ DB2 V8.1

## **Employee Online application**

The main purpose of the Employee Online application is to serve the ITSO-Electronics.com HR department and all employees. It has well-established periodic requirement peaks, that is, at the end of each month and at the end of the financial year.

Since Employee Online is solely for ITSO-Electronics.com internal use, very low usage is recorded after office hours and during weekends.

The components that make up the Employee Online application are the following:

### ***Internal WebSEAL***

Internal WebSEAL is the security component that acts as the front gate to provide secure Web access to Employee Online (refer to “Implementation of IBM Identity Management solution” on page 183 for more details about WebSEAL). It runs with the following software:

- ▶ Linux Red Hat 7.3
- ▶ IBM Tivoli Access Manager for e-business WebSEAL 4.1

### ***Employee Online Web and application server***

The Web and application server is running the following software:

- ▶ AIX 5.1
- ▶ Apache HTTP Server 1.3
- ▶ IBM WebSphere Application Server 5.0.2
- ▶ Employee Online IBM WebSphere Application Server component

### ***Stocks Online database***

The Employee Online database server runs the following software:

- ▶ z/OS V1R4.0
- ▶ DB2 V8.1

## **16.4.2 Business objectives**

The CIO of ITSO-Electronics.com has defined the following requirements to meet the challenges described previously:

- ▶ The Stocks Online Web application infrastructure should be always available and ready to address consumer demands, 24 hours a day, 7 days a week, all year round.
- ▶ The Stocks Online Web application infrastructure should be optimized, and maximum utilization obtained for new and existing investments.
- ▶ The Stocks Online Web application infrastructure can quickly recover from any failure, taking minutes instead of hours or days.
- ▶ Peaks in demand are met quickly and effectively and without business disruption.
- ▶ The infrastructure should be scalable, allowing it to grow and shrink with changing demands without sacrificing manageability of complex, heterogeneous environments.

- ▶ The integrity and privacy of customer information must be maintained. Transactions need to be executed in an environment that customers feel is safe and secure.

### 16.4.3 Technical objectives

ITSO-Electronics.com wants to get more from their existing IT infrastructures. They need to optimize manually-intensive processes and environments to adjust for rapidly-changing business conditions, as well as keep pace with dynamic and often fluctuating demands of key business applications, without incurring overly-expensive capital costs. Let us summarize the difficulties and concerns that ITSO-Electronics.com is facing today:

- ▶ Difficulties in managing and increasing service levels:
  - Without effective tools to manage the complex heterogeneous environment, more personnel are required to manage the infrastructure to ensure service level agreements (SLAs) are met.
  - In addition, reporting on SLA attainment is a time-consuming and labor-intensive process. It requires manual comparison between data from multiple IT infrastructure components, and the customer-agreed SLA.
  - Traditional system management tools only focus on servers, routers, applications, disk space, and so on. These types of reports do not show which business aspects are affected if any of the components is down or unavailable.
- ▶ Uneven resource allocation and utilization results in high IT costs:

As described in 16.4.1, “Current environment” on page 219, Stocks Online needs additional resources to meet the demand from users, while Employees Online may have enough or even extra resources allocated to it.
- ▶ Insufficient infrastructure and resources to respond to business changes quickly and with flexibility.
- ▶ Increasing amounts of risk due to inefficient resource management.

### 16.4.4 Solution approach

IBM Web Infrastructure Orchestration is one of the key components of the IBM Infrastructure Management family of offerings. By proactively sensing and responding to peaks in demand, then re-allocating IT resources to the most critical business processes based on ITSO-Electronics.com’s specific business policies, it helps ITSO-Electronics.com to utilize the resources allocated to processes (such as Stocks Online) where they support a specific customer’s business requirements most efficiently.



The IBM Tivoli Intelligent ThinkDynamic Orchestrator offers a powerful system that can:

- ▶ Gather information about the performance of all ITSO-Electronics.com application clusters and build a workload model that can predict resource requirements going forward
- ▶ Manage resources across all ITSO-Electronics.com application clusters to optimize business-aligned service level delivery
- ▶ Automate the deployment of the optimal computing resources to each application environment

IBM Web Infrastructure Orchestration is a pre-tested, pre-integrated solution of software and the IBM @server BladeCenter along with optional services and packaged intelligence that automatically modifies or adds resources to client's Web-serving environments. It adjusts Web capacity immediately, streamlining operations, optimizing resources, and lowering management costs.

Key software components include:

- ▶ IBM Tivoli Intelligent ThinkDynamic Orchestrator
- ▶ IBM Tivoli Provisioning Manager
- ▶ IBM Tivoli Monitoring
- ▶ IBM Tivoli Monitoring for Web Infrastructure
- ▶ IBM Tivoli Monitoring for Databases
- ▶ IBM Tivoli Storage Manager Extended Edition
- ▶ IBM Tivoli Storage Manager for Application Servers
- ▶ IBM WebSphere Application Server Network Deployment
- ▶ IBM DB2 UDB Workgroup Server
- ▶ IBM Director
- ▶ IBM Remote Deployment Manager

Key hardware components include:

- ▶ IBM @server Blade Center
- ▶ IBM @server xSeries 360 Server
- ▶ IBM FASTT900 Storage Server

The solution also includes packaged intelligence, or policies; user-defined specifications of service levels and priorities; and workflows. These workflows are sequences of automated administration and management steps, as well as IBM Director scripts designed specifically for this environment. The entire system is highly extensible, enabling clients to scale their infrastructure on a pay-as-they-go basis.

## Implementation of IBM Web Infrastructure Orchestration

The decision varies from one company to another on how many components of the offering should be deployed. In ITSO-Electronics.com's case, after several discussion with the CIO, the solution designers identified three key areas to be addressed:

- ▶ Provisioning infrastructure
- ▶ Managing capacity
- ▶ Delivering service level

For the ITSO-Electronics.com scenario, we used IBM Tivoli Intelligent ThinkDynamic Orchestrator as the starting point. This combines autonomic capabilities and orchestration technology to automatically shift resources such as servers, storage, and bandwidth as required. It can sense shifts in demand based on defined business rules and automatically take action to re-allocate them accordingly, provisioning resources throughout the IT infrastructure. Thus, ITSO-Electronics.com can gain access to under-used computing systems without having to invest in additional capacity.

Figure 16-2 shows the orchestrated provisioning in the ITSO-Electronics.com San Diego data center. IBM Tivoli Intelligent ThinkDynamic Orchestrator operates in a closed loop that performs automatic resource requirements prediction, based on pre-defined service level objectives and agreements, and automates infrastructure deployment. This cycle ensures each application has the resources it needs, when it needs them, without static over provisioning.

Since any application can be allocated on any server at any time, a server could be built specifically for an application when it needs it and is subsequently de-allocated when it is no longer needed by the application.

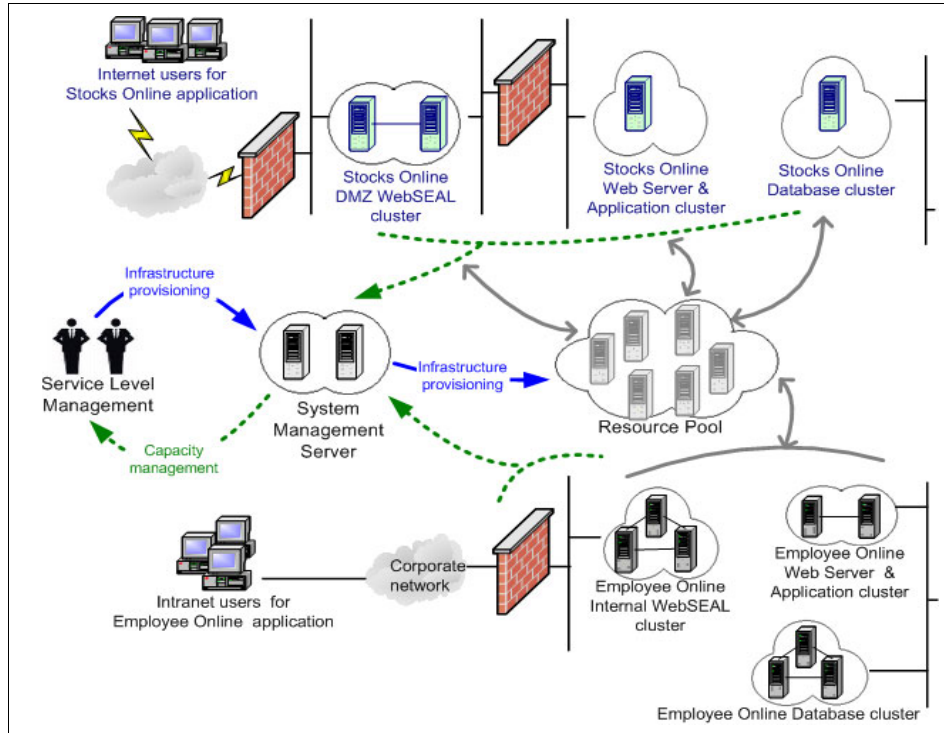


Figure 16-2 Orchestrated provisioning in San Diego data center

Each application consists of a clustered WebSEAL component, a clustered Web and application server, and a clustered database back-end component.

## Clusters

There are three main components within a cluster:

- ▶ **Cluster Manager:** The cluster manager, also known as a load balancer, distributes workload to the participating members of the cluster or replicas. Users of the cluster know the cluster via a virtual IP address (VIP). The cluster manager/load balancer can distribute traffic to replicas based on various techniques, such as round robin and least busy algorithm.
- ▶ **Replicas:** All the replicas within the cluster are identical. They are the servers that serve requests. Workload is distributed to these systems via the load balancer. The parallel nature of replicated servers and services often leads to both improved scalability and increased availability.

- ▶ **Reserve systems:** This component is implied in a cluster. Clusters use a horizontal scaling technique to grow and possibly shrink. Horizontal scaling is a common technique in the Web space and is simply adding or subtracting instances of the same service as the workload changes.

The basic tasks that are required to manage the cluster include:

- ▶ **Information gathering:** Gathering real-time data such as the number of requests and resource consumption on the replicas directly from the load balancers and the replicas via proxy.
- ▶ **Decision making:** Determining which resource should be used based on application priority and control the allocation of resources.
- ▶ **Cluster visualizing/configuration:** Using the ability to see, update, and enhance the cluster as a whole.
- ▶ **Automated provisioning:** Automating the provisioning and de-provisioning of system resources in an unattended, controllable, verifiable, and automated operation or workflow.

## IBM Web Infrastructure Orchestration components

Figure 16-3 on page 227 shows a high-level view of the IBM Tivoli Intelligent ThinkDynamic Orchestrator architecture. There are six main components: data acquisition engine, application controller, deployment engine, data center model, global resource manager, and management interface. A brief description of each component follows. For more details, refer to the *IBM Tivoli Intelligent Orchestrator Overview Guide*, SC32-1419.

- ▶ **Data Acquisition Engine:** The data acquisition engine is responsible for the information gathering task. It acquires and pre-processes performance metric data from each managed application environment. It also distributes signals to other components of the IBM Tivoli Intelligent ThinkDynamic Orchestrator.
- ▶ **Application Controller:** The application controller participates in the decision making task as an instance running on each application environment under management. Based on the workload model and predictions, as well as on real-time performance data acquired from the data acquisition engine, it then determines the resource requirements of the application.
- ▶ **Global Resource Manager:** The global resource manager also participates in the decision making task. It receives requirements for servers or network devices from all application controllers, and manages the overall optimization of data center assets. It has two primary responsibilities:
  1. Make optimal resource allocation decisions.
  2. Ensure stable control over the application infrastructure.

- ▶ **Data Center Model:** The data center model participates in the cluster visualizing/configuring task. It represents all of the physical and logical assets under IBM Tivoli Intelligent ThinkDynamic Orchestrator's management, such as servers, switches, load balancers, application software, VLANs, security policies, service level agreements, and so on. It keeps track of the data center's assets and associated allocations to customer applications.
- ▶ **Management Interface:** The management interface is responsible for the cluster visualizing/configuring task. It provides an overview of the state of all physical and logical assets in the data center infrastructure, offering information about the servers and their allocation, and generating configurations and allocations. It can also be used to create application environments.
- ▶ **Deployment Engine:** The deployment engine participates in the automated provisioning task. It is responsible for the creation, storage, and execution of repeatable workflows that automate the configuration and allocation of assets.

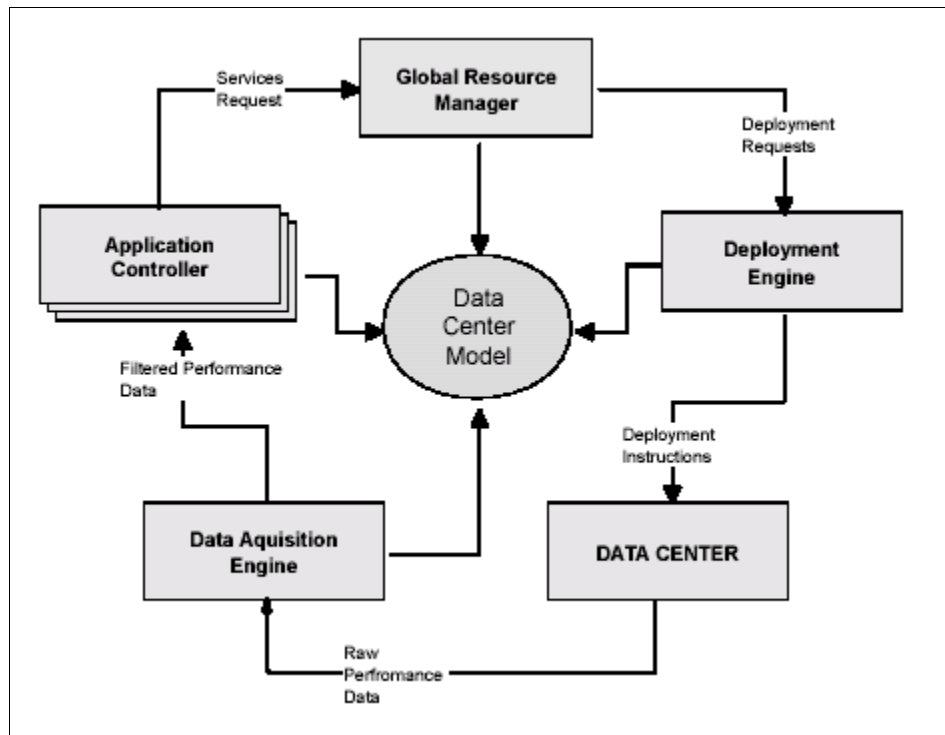


Figure 16-3 IBM Tivoli Intelligent ThinkDynamic Orchestrator architecture

Once the environment is deployed, the process is executed as follows:

1. The data acquisition engine collects data from devices.
2. The data acquisition engine sends data to the application controller, which evaluates server needs.
3. The application controller sends its needs to the global resource manager.
4. The recommendation broker weighs different requests from application controllers and generates recommendations that optimize the health of the data center as a whole based on preset SLA values.
5. The deployment engine translates recommendations into workflows.
6. The deployment engine executes workflows on managed servers.

ITSO-Electronics.com's Web infrastructure includes the Apache HTTP Server, the IBM WebSphere Application Server, the Stocks Online IBM WebSphere Application Server component, the Employee Online IBM WebSphere Application Server component, DB2, and the management agents in the blade servers within the IBM eServer BladeCenter.

The software stack that will be deployed on a blade server is captured as an installation image that will be installed by the IBM Remote Deployment Manager tool in IBM Director that is offered for xSeries systems. The installation image includes the operating system, management agents, and the middleware components. Installation images will be resident on the local storage of the domain manager.

The components that make up the solution for ITSO-Electronics.com are shown in Figure 16-4.

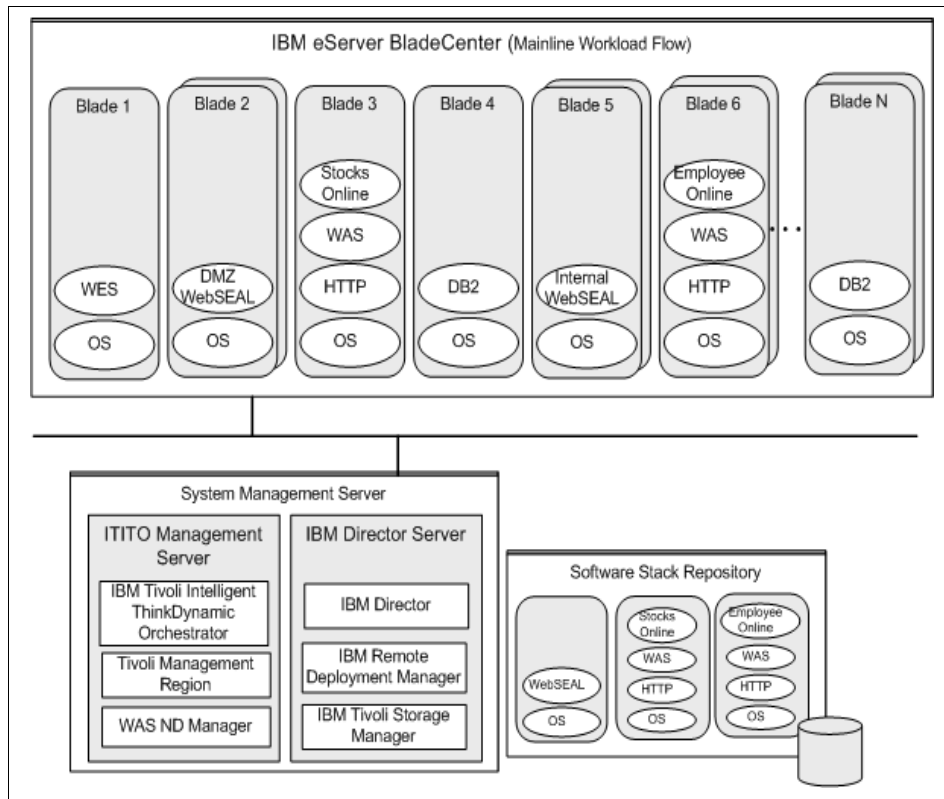


Figure 16-4 IBM Web Infrastructure Orchestration for ITSO-Electronics.com

### System Management server

Figure 16-4 shows an overview of the IBM Web Infrastructure Orchestration components that are used in ITSO-Electronics.com. Let's take a closer look at the system management server component.

The system management server is made up of several management tools including IBM Tivoli Monitor, IBM Director, Remote Deployment Manager, and others. The IBM Tivoli Intelligent ThinkDynamic Orchestrator acts as the single point of control to coordinate all the separate management tools within the management domain.

With object model and abstract logical device operations, IBM Tivoli Intelligent ThinkDynamic Orchestrator can utilize different implementations for different devices and management tools. It leverages other tools by delegating specific operations and coordinating the results through its workflow engine.

The IBM Web Infrastructure Orchestration solution architecture uses the delegation model, integrating multiple management products which serve specific roles. IBM Director, RDM, WebSphere Application Server Network Deployment, and the IBM Tivoli Monitor server all provide different management domains within the offering. The roles of the management tools are summarized as follows:

- ▶ **IBM Tivoli Intelligent ThinkDynamic Orchestrator:** Overall orchestration of provisioning workflows
- ▶ **IBM Director:** Hardware control, fault monitoring, blade insertion/removal detection
- ▶ **IBM Remote Deployment Manager:** Bare-metal operating system and agent installation, hardware configuration, firmware updates
- ▶ **WebSphere Application Server Network Dispatcher:** WebSphere Application Server cluster management and HTTP server configuration
- ▶ **Tivoli Management Region server:** Hosts the IBM Tivoli Monitoring application which provides application-specific monitoring for the HTTP, WebSphere Application Server, and DB2 tiers

The IBM Director, Remote Deployment Manager, WebSphere Application Server Network Dispatcher, and Tivoli Management Region servers serve as management proxies. All the management applications can be installed on one physical management server or put into separate management servers as required.

## Setting up the environment

Now, let's discuss the tasks and processes that IBM Tivoli Intelligent ThinkDynamic Orchestrator goes through to implement ITSO-Electronics.com's solution.

### Task 1: Information gathering

To start the process, raw data is gathered and turn it into useful information. Table 16-1 shows the customer data that ITSO-Electronics.com gathered.

Table 16-1 ITSO-Electronics.com customer data

Line of business	Application / description	Service level objective	Network traffic utilization
Retail	Stocks Online - Online Web application for ITSO-Electronics.com customers for stocks and funds inquiry and transactions	Highest priority	Peaks during stock market operating hours; higher spikes at month end
Human Resources	Employees Online - Internal employee information system	Medium to High priority	High utilization during last 4 days of every month



Infrastructure assets are both physical and logical. They include physical network entities such as load balancers, switches, routers, and firewalls; as well as VLANs, subnet, protocols, and so on. Once defined, IBM Tivoli Intelligent ThinkDynamic Orchestrator will treat all assets as data objects.

Information such as boot servers, terminal servers, and power units may also need to be collected depending upon the structure of the ITSO-Electronics.com data center. Figure 16-5 shows the logical network diagram of the ITSO-Electronics.com San Diego data center. The diagram is then transformed into logical network data.

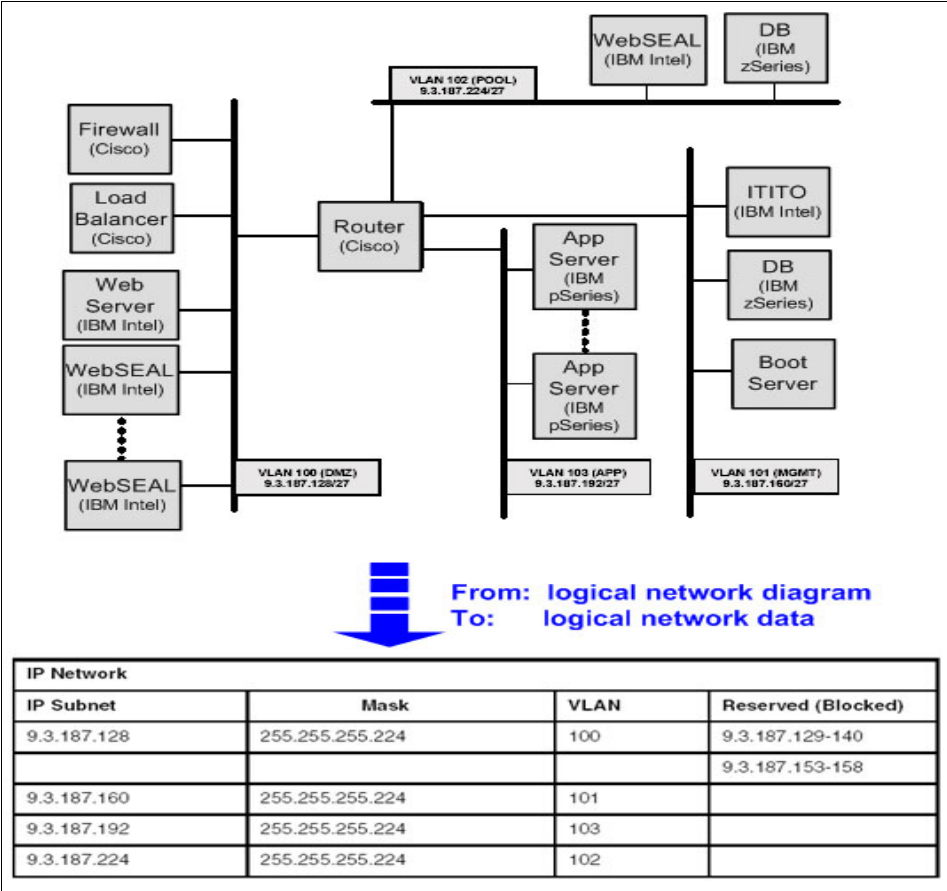


Figure 16-5 Logical network data of ITSO-Electronics.com San Diego data center

Repeat the process to gather other information such as physical network switches, physical servers on network, and load balancers, and transform it into useful logical data.

## ***Task 2: Decision making***

After collecting the infrastructure data from Task 1, we can analyze the data for commonality and for completeness. From this point on in the planning cycle, we will be creating “derived data” in IBM Tivoli Intelligent ThinkDynamic Orchestrator terms. That is, we will be making decisions about how we will set up provisioning based on this collected information, using the following processes and terms:

- ▶ **Cluster Organizations:** To define the groups of server “Clusters” that run with the same service.
- ▶ **Resource Pools:** To define the groups of physical or logical servers that are available to be provisioned “on demand”. The servers in the pool can be moved in or out of the pool based on the needs of the application clusters.
- ▶ **Provisioning:** To define the minimum and maximum number of servers each customer could have provisioned in their application cluster at any one time, to determine the minimum and maximum response time within the application, and the minimum percentage of time available.
- ▶ **Workflows:** To define the workflow for a resource that needs to be added or removed either manually or automatically from a resource pool to an application cluster.
- ▶ **Software Stacks:** To define the list of software products in sequential order. There are two type of stacks, image and product. The former is the stack’s initial load served from a boot server with an optional ordered list of product stacks, products, and patches; the latter is an ordered list of products stacks, products, and patches.
- ▶ **Service Access Point:** To determine the protocols and credentials that identifies which network protocols ITSO-Electronics.com will use when communicating with devices on the network. Some of the common access points are SNMP, secure shell, telnet, ftp, icmp, and so forth.
- ▶ **Networks:** It is not necessary to define all devices on the network, but it is important to reserve some IP addresses for future use, or those IP addresses that are currently in use but are not managed by IBM Tivoli Intelligent ThinkDynamic Orchestrator.

## ***Task 3: Cluster visualization and configuration***

With the plans and materials from Tasks 1 and 2, we can now start to build the data center model for ITSO-Electronics.com.

### ***Data center model***

The data center model represents the physical and logical assets under IBM Tivoli Intelligent ThinkDynamic Orchestrator’s management, such as servers, switches, load balancers, application software, and so on. It keeps track of the data center’s assets and associated allocations to customer applications.

We need to configure the blade servers, boot servers, terminal servers, access control lists, firewalls, load balancers, power units, routers, servers, switch fabrics, subnets, license pools, software products, software patches, and software stacks — and group them all under the data center assets and resources menu — while the applications and clusters are grouped under the customer applications.

### ***Data center model asset relationship***

The task here is to link all the related components together. For example, a subnet is linked with a switch-fabric. Figure 16-6 shows the high-level assets relationship of Stocks Online and Employee Online applications in ITSO-Electronics.com.

**Note:** The grey boxes are the starting points.

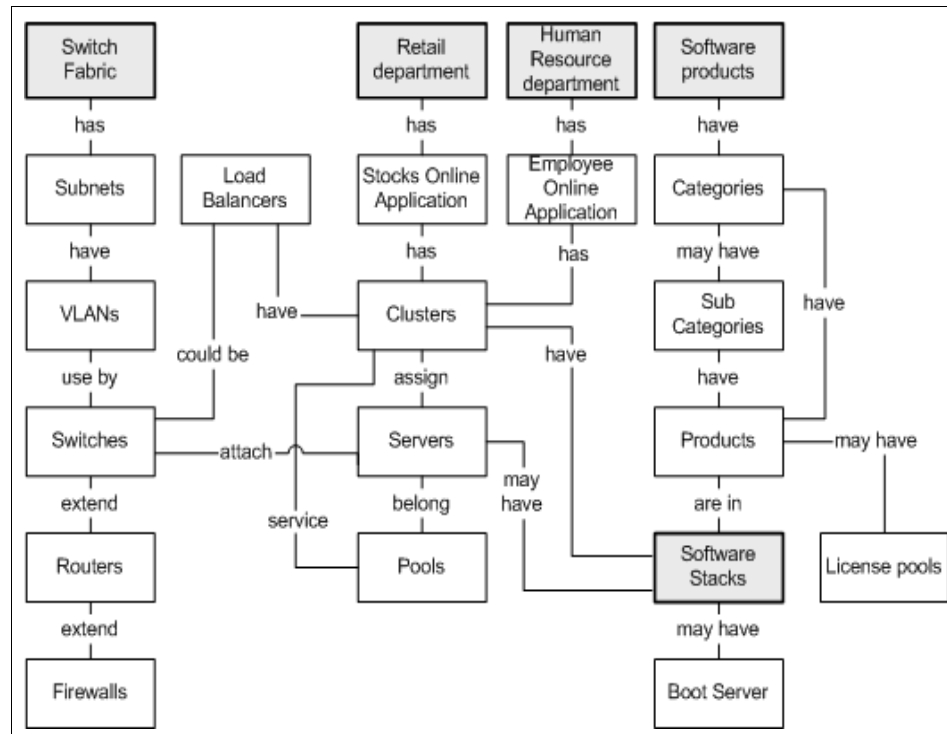


Figure 16-6 High-level asset relationship in ITSO-Electronics.com

### ***Building a data center model***

With all the necessary information described in the previous sections, we can start building the data center model for ITSO-Electronics.com via the Web interface or the command line (XML) methods that are provided by IBM Tivoli Intelligent ThinkDynamic Orchestrator.

### ***Workflow***

In order to automate the processes, there are a few more steps that need to be performed so that the processes know how to deploy the defined assets:

- ▶ Create workflows to automate the processes.
- ▶ Create automation packages so that we can reuse the workflows.
- ▶ Update the data center model as required.

Once the data center model goes “live” and the mode is automatic, updates to devices in the model can occur frequently. At this point, the startup XML is no longer an accurate representation of the data center.

Additions to the data center model (that is, new devices) can and will likely be made via XML. Minor changes to devices will most likely be made via the GUI.

### ***Task 4: Automated provisioning***

A deployment is the actual provisioning of servers in the data center environment. The deployment engine is responsible for this task. It includes the creation, the storage, and the execution of repeatable workflows that automate the server configuration and allocation in the system.

Deployments can happen manually through IBM Tivoli Intelligent ThinkDynamic Orchestrator management interface or automatically based on a recommendation from the resource broker. During a deployment, a workflow is executed and servers and network devices in the environment are changed. For example, IP addresses change, servers are added to load balancers, software is installed, and so on.

### **Provisioning scenarios**

Stocks Online is a business-critical application. During weekdays, especially during Stock Market operating hours, there is a predictable increase in the number of user and transaction requests. If the current system resources are under-allocated, ITSO-Electronics.com receives customer complaints about the poor response of the Stocks Online system and that it causes them to suffer from not reacting to the market changes in a reasonable time.

### ***Addition or removal of system resources for Stocks Online***

The data acquisition engine keeps acquiring and pre-processing performance data from the Stocks Online application environment. The performance data is captured from the Stocks Online application, the operating system, the network, and the infrastructure layers. The captured raw data is then filtered and sent to the application controller of the Stocks Online application.

Based on Stocks Online's workload model and predictions, and using the real-time performance data that is received from the data acquisition engine, the Stocks Online application controller determines Stocks Online is experiencing increased workload that it is over the pre-defined service level threshold. The application controller suggests additional system resources are required for Stocks Online, which includes DMZ WebSEAL, the Stocks Online Web and application server, and the back-end database. It then sends its requirements to the global resource manager.

When the global resource manager receives the requirements for servers from the Stocks Online Application controller, it makes the optimal server allocation decision. It also has to ensure stable control over the Stocks Online application infrastructure. After considering the different server requirements for Stocks Online, the server reconfiguration and allocation data is passed on to the deployment engine.

When the deployment engine receives the recommended deployment request from the global resource manager, it immediately translates it into workflows and executes the workflows on its managed servers. The workflow includes:

- ▶ Deployment of operating system, middleware, and software to blades.
- ▶ Configuration of operating system, middleware, software on a blade, and dynamic inclusion of the blade into the Stocks Online application cluster.

As a result, three additional servers (one for DMZ WebSEAL, one for Stocks Online Web and application, and one for the back-end database), each with proper setup and configuration, are added from the resource pool to the Stocks Online application. An overview of the process is shown in Figure 16-7.

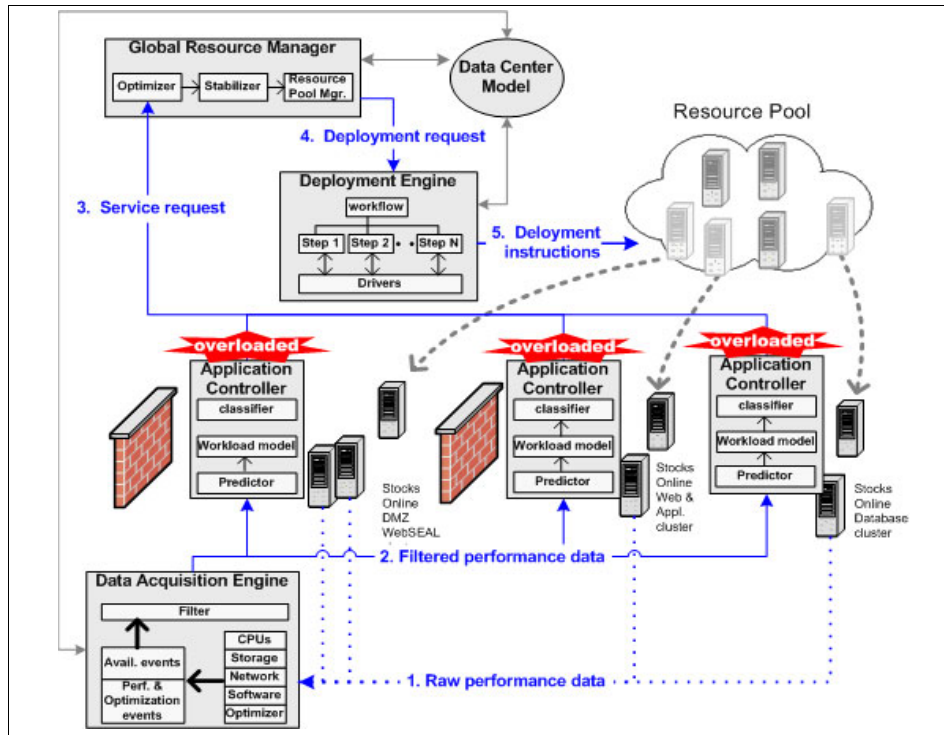


Figure 16-7 Additional resource is required for Stocks Online application

The data acquisition engine continues to gather and monitor data from the devices. Should the workload of the Stocks Online application drop to a pre-defined level during weekend or holidays, the whole cycle will be restarted, but this time, the recommendation from the global resource manager is to de-provision servers from the Stocks Online application as the resources are under-utilized.

Figure 16-8 shows that the Internet traffic has dropped to a pre-defined “Normal” level during the weekend. The data acquisition engine filters the captured performance data from its managed devices and sends it to the application controller.

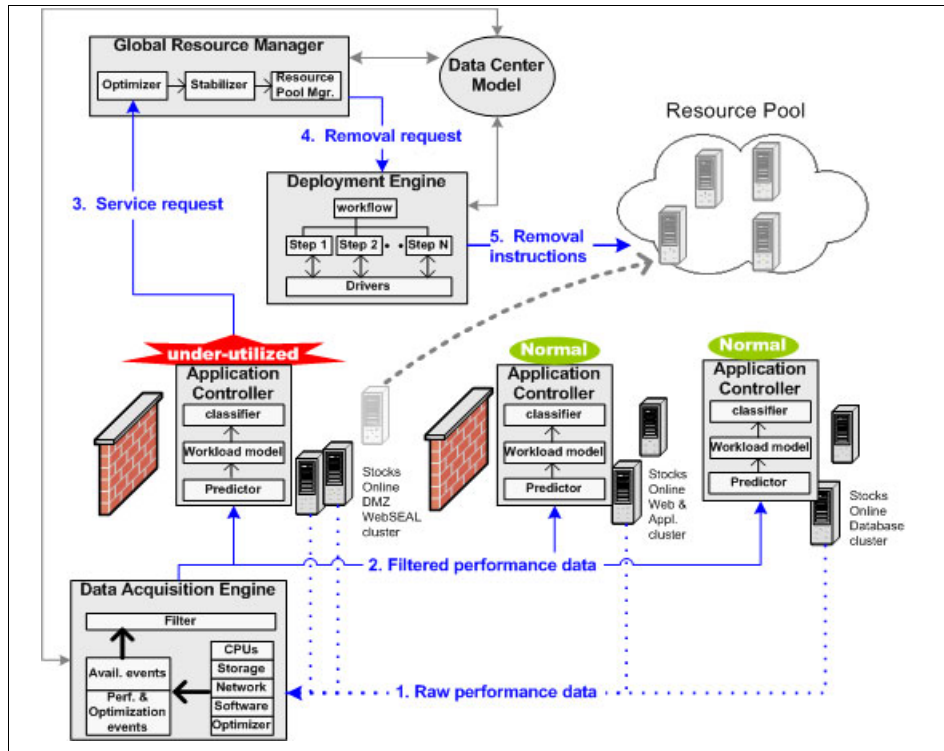


Figure 16-8 Removal of resource for Stocks Online application

With the workload model and predictions, and using the real-time performance data received from the data acquisition engine, the application controller determines the additional DMZ WebSEAL server is no longer required to serve the Internet traffic since there are already two WebSEAL servers in place. However, the back-end Web and application server, together with the database server, are still required. With the additional servers that were previously assigned to these two clusters, the workload can be handled. The application controller suggests to remove one of the DMZ WebSEAL servers and return it to the resource pool. It then sends this request to the global resource manager.

The global resource manager receives the recommendation from the application controller and makes the optimal server allocation decision. It also ensures the Stocks Online application infrastructure is in a stable control environment. The DMZ WebSEAL server unconfiguration and de-allocation data is passed on to the deployment engine.

The deployment engine follows the pre-defined workflows and procedures to execute the following actions:

- De-provisioning of operating system, middleware, and software from the DMZ WebSEAL server.

Now, redundant system resources (such as the additional DMZ WebSEAL server) are removed from the Stocks Online application and returned back to the resource pool.

### ***Additional Web components of Employee Online***

In this scenario, we describe how a Web component (HTTP and WebSphere Application Server tier) is provisioned on demand.

The Employee Online application has a record of under 10% system resource utilization after office hours, during public holidays, and on weekends. However, at month-end and year-end, both the HR department and employees experience slow response from the Employee Online application and sometimes, employees are prevented from accessing the system due to server overloading.

As shown in Figure 16-9, Employee Online is a highly sensitive Internal application that has four tiers:

1. The front-end load balancer that distributes the load to the internal WebSEAL servers.
2. The Internal WebSEAL servers that serve as the front-gate (secure reverse Web proxy) to authenticate and authorize users before entering the Employee Online application.
3. The Web and application server that provides the Web-based user interface and contains business logic.
4. The back-end database server that contains all the employee-related information.

During regular data acquisition from the data acquisition engine, the Employee Online WebSphere application server workload exceeding the pre-defined performance threshold is detected. Based on the policy-based analysis, the application controller determines it is a service level agreement violation. It decides additional system resource is required to be allocated to the Employee Online Web and application cluster.

The request is sent to the global resource manager, which then makes the optimal system resource allocation decision. It also has to ensure stable control over the Employee Online application infrastructure. After considering the different server requirements for Employee Online, the server reconfiguration and allocation data is passed on to the deployment engine.



After the deployment engine receives the deployment request from the global resource manager, it translates it into these workflow execution commands:

1. Assign server static IP addresses on each network interface card on the appropriate VLAN.
2. Update the associated switch ports to the correct VLANs.
3. Restart the Tivoli Monitor Agent on the endpoint to notify the Tivoli Management Region gateway of the IP address change.
4. Install the software stack associated with the cluster. This operation installs the HTTP Server and the WebSphere Application Server and Employee Online application.
5. Update the configuration and start the server.
6. Download and activate the correct IBM Tivoli Monitor resource monitors on the server.
7. Update the load balancer configuration.

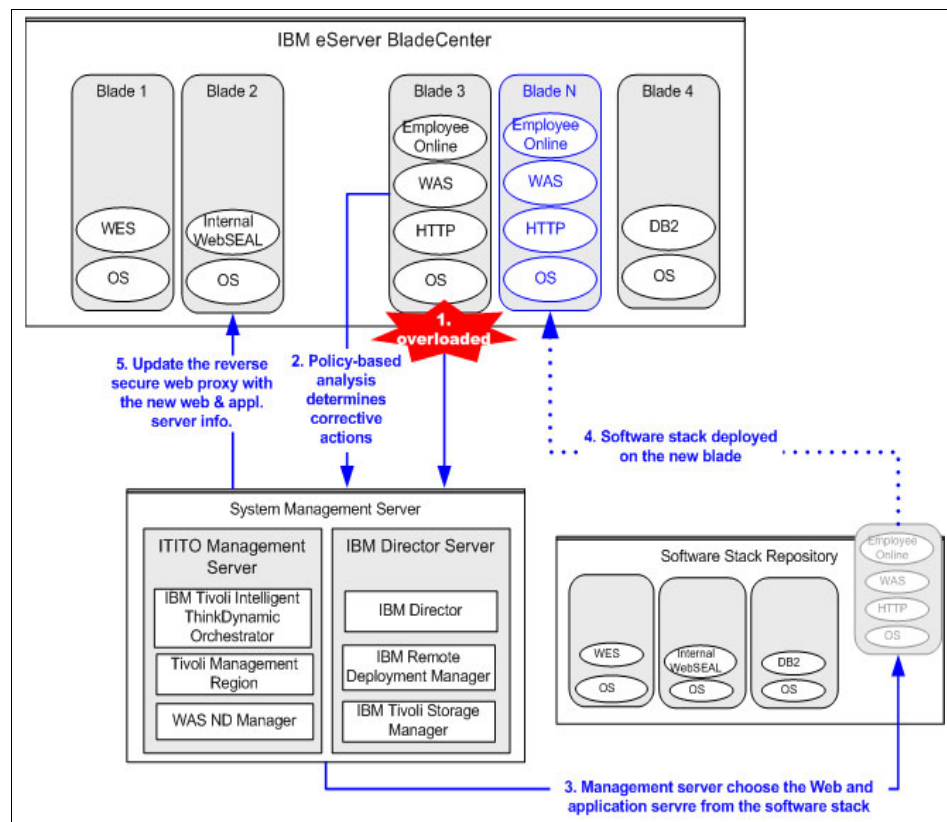


Figure 16-9 Provisioning of server to the Employee Online application

### ***System resource failure***

In the case of the IBM Director Management Processor Assistant detecting a hardware failure or a predicted failure on a blade server, an event action plan will initiate the WIP Blade Failed workflow on the IBM Tivoli Intelligent ThinkDynamic Orchestrator server. This workflow will automatically perform the following actions:

1. Transition the server into the failed state.
2. Force a removal of the failed server from the active cluster.

## **16.4.5 Benefits and summary**

Once fully deployed, IBM's Infrastructure Management solution will form an active linkage between Business Service Management and the automation disciplines of availability, security, optimization and provisioning. This linkage is maintained by a perpetual cycle of real-time performance and event data analysis, enforcement of customer-defined business policies, and automated change to the underlying infrastructure.

Based on ITSO-Electronics.com's defined policies for a collection of e-business applications, the orchestrator will automatically assign computing resources where and when they are needed to maintain application service levels without costly over-provisioning. By leveraging real-time data and customer-defined business policies based on business or financial priorities to take intelligent, automated action, orchestration delivers increased utilization of existing and new resources, improved productivity of IT staff, accelerated responsiveness to changing business needs, and elevated application service levels.

ITSO-Electronics.com now can enjoy the following benefits:

- ▶ Automated Orchestration, increasing utilization directly tied to business results by proactively sensing and responding to peaks in demand and allocating IT infrastructure using business policies.
- ▶ ITSO-Electronics.com can manipulate the IT environment in real time according to defined business policies to achieve desired business goals.
- ▶ Orchestration does this by sensing increases in demand for resources and automatically taking action to allocate / re-allocate resources accordingly, and by provisioning resources throughout the entire system — hardware, software, and applications.
- ▶ By dynamically allocating capacity to applications that require it, utilization of under-utilized computing systems is improved without requiring an investment in additional capacity, and the company's systems "sense and respond" to disruption or threats before they occur.

## 16.5 Product positioning

We have seen that solving the provisioning issues impacts many parts of the general On Demand Operating Environment. To make this kind of environment a reality requires software as well as hardware capabilities. In the following sections we describe some of the related products and features that are available to start building this kind of environment today.

### 16.5.1 Self-configuring IBM @server

There are many features that are integral to IBM hardware and storage platforms which can be used to provision additional hardware resources. Capacity on demand features (permanent and/or temporary) in servers and storage allow a company to quickly acquire additional capacity for hardware resources, often without having to experience any down time. Workload management capabilities and intelligent resource director capabilities in the hardware allow those resources to be dynamically reconfigured and allocated to workload based on business goals.

IBM led the industry in introducing IBM @server On/Off Capacity on Demand (On/Off CoD), first available on its mainframes in 1998, and subsequently extended to its other product lines. Currently, select IBM @server and TotalStorage products offer On/off CoD, designed to help provide extra processing power or storage capacity to meet the demands of e-business without affecting current operational commitments. By listening to and working with its customers, IBM has developed a series of On/Off CoD services that can help address a full range of business needs.

Many on demand offerings are available from IBM for a number of models across the @server and TotalStorage product lines. @server examples include: IBM @server Capacity Upgrade on Demand (CUoD), IBM @server On/Off Capacity on Demand (On/Off CoD), Capacity BackUp, and others. TotalStorage examples include: Standby Capacity on Demand (SCoD) for disk and tape and connectivity on demand (SAN).

CUoD and SCoD provide planned growth for customers who know they will need expanded processing capabilities in the future. Standby capacity may be incorporated into the original purchase, but some up-front charges, pricing premiums, or purchase commitments may be required for this additional capacity. Customers pay to activate the capacity when it is needed, just as they would for traditional upgrades, but CUoD avoids the purchasing, approval, shipment, and installation processes prior to placing the additional resources online. Depending upon the specific product, IBM offers CUoD for processors, memory, I/O channels, blade servers, and disk storage.

On/Off CoD provides temporary increases in processor capacity for the peaks and valleys that occur in any business. On/Off CoD differs from CUoD and Standby CUoD in that its capacity can be activated from a customer's original configuration or deactivated as business demands dictate. The specific resources available for On/Off CoD differ by IBM @server product and are sold in prepaid time-period increments or monitored usage (pay as you go).

Logical Partitioning (LPAR) is the allocation of system resources to create logically separate systems within the same physical footprint. System Resources on iSeries include:

- ▶ Processors
- ▶ Main Storage
- ▶ Interactive performance
- ▶ System buses
- ▶ Disk
- ▶ I/O controllers and devices

With the click of a mouse, under a schedule, or under program control, one can move resources dynamically to a partition to alleviate a resource problem, or just to prepare for different usage patterns. Capacity on demand further enhances this functionality with resources packaged with the system that can be activated on an “as-needed” basis (and paid for as such).

### **IBM @server zSeries**

This brand includes a wide variety of offerings related to on demand. If we start with permanent activation of resources, there is the possibility to increase the number of processors on a system (available for z800, z900 and z990) and/or memory (available for z900 and z990) up to the maximum installed. With the Express option on Customer Initiated Upgrade (CIU), an upgrade can be installed within a few hours after order submission.

On/Off Capacity on Demand (On/Off CoD) is designed to deliver short-term additional capacity. On/Off CoD is designed to temporarily turn on as Central Processors (CPs), previously uncharacterized Processor Units (PUs), unassigned CPs, and unassigned Integrated Facilities for Linux (IFLs) that are available within the current model. On/Off CoD is delivered through the function of Customer Initiated Upgrade (CIU). To participate in this offering, one must have installed CIU Enablement (FC 9898), and On/Off CoD Enablement (FC 9896). Subsequently, one may concurrently install temporary capacity by ordering On/Off CoD Active CP (FC 9897) up to the number of current CPs, and use it for an indeterminate time. One is then billed on a twenty-four hour basis.

In addition, IBM zSeries servers offer Capacity BackUp offerings, where dedicated processors can be activated when a major disaster has stopped a remote system.

z/VM offers unique functions and technology on the zSeries that take advantage of the virtualization capabilities. z/VM gives the zSeries server the ability to virtually partition each Logical Partition (LPAR) into many virtual servers, to virtualize processor, communication, memory, storage, I/O and networking resources. In particular, z/VM offers the ability to host, on a single zSeries server tens to hundreds of Linux images, allowing for the consolidation of UNIX, Microsoft Windows, and Linux workloads on a single physical box. The Linux servers then share the hardware resources. z/VM SCSI support, available since z/VM V5.1, allows a Linux server farm to be deployed on z/VM in a configuration that includes only SCSI disks (ECKD™ disks are no longer required).

This is IBM intent, at the time of writing, to offer the IBM Virtualization Engine support for Linux on zSeries; that would include the support of Enterprise Workload Manager, IBM Director Multiplatform and the IBM Dynamic Infrastructure for mySAP Business Suite.

### **IBM @server pSeries**

The pSeries brand carries similar offerings. The permanent activation of resources includes processors and memory on high and mid range systems (p650, p670 and p690). The minimum increment is two processors and/or 4 GB of memory, up to the maximum installed on the system. The On/Off CoD offering is implemented a bit differently. A feature code which gives access to 30 days of activation can be purchased up front. The processors can then be activated any time they are required for a minimum duration of 24 hours.

In October 2003, IBM announced a new offering in the Capacity BackUp (CBU) area. It is now possible to order a system with CBU processors. These processors will be activated in case of a major disaster. It is also possible to use those processors with On/Off CoD.

The pSeries 690 server is the first pSeries server to incorporate the ability of being partitioned. Its architectural design brings logical partitioning to the UNIX world, capable of multiple partitions inside a single server, with great flexibility in resource selection. The partitioning implementation on pSeries 690 differs from those of other UNIX system vendors in that the physical resources that can be assigned to a partition are not limited by internal physical system board boundaries.

Processors, memory, and I/O slots can be allocated to any partition, regardless of their locality. For example, two processors on the same POWER4™ silicon chip can be in different partitions. Peripheral component interconnect (PCI) slots are assigned individually to partitions, and memory can be allocated in fixed-size increments. The fine granularity of the resources that can be assigned to partitions provides flexibility to create systems with the desired resources.

The partitioning-capable pSeries servers are also capable of running both AIX 5L Version 5.1 or later and Linux inside a partition on the single system simultaneously.

In 2004, the pSeries p5 systems introduced new virtualization system technologies; by automatically applying only the required amount of resource needed for each partition, they allow an increased overall utilization of systems resources. They include the following technologies:

- ▶ The POWER Hypervisor, which supports partitioning and dynamic resource movement across multiple operating systems environments
- ▶ The micro-partitioning, which enables a logical partition to use less than a full physical processor
- ▶ The Virtual Ethernet, which enables inter-partition communication without the need for physical network adapters assigned to each partition
- ▶ The Virtual I/O, which provides the ability to dedicate I/O adapters and devices to a virtual server; shared devices and adapters can be used simultaneously by more than one partition

For more information, please refer to the following redbooks: *Advanced POWER Virtualization on IBM eServer p5 Servers: Introduction and basic Configuration*, SG24-7940, and *Introduction to pSeries Provisioning*, SG24-6389.

## **IBM @server iSeries**

Every iSeries Model 825, 870, or 890 comes with extra processor capacity built into the server. This extra capacity, delivered as stand-by processors, can be activated permanently or temporarily.

IBM @server Capacity Upgrade on Demand for iSeries enables the growth of processor capacity as requirements grow. Stand-by processors can be permanently activated by ordering an activation feature that results in an activation code being provided that can quickly be entered at the server console.

IBM @server On/Off Capacity on Demand for iSeries delivers great flexibility in meeting the peak demands of a dynamic business. Stand-by processors can be temporarily activated with a simple request right at the server console, provided the server has been enabled for use of temporary capacity.

The iSeries for Capacity BackUp offering is intended for companies requiring an off-site, disaster recovery machine at an extremely affordable price. Using iSeries On/Off CoD capabilities, the iSeries for Capacity BackUp offering has a minimum set of startup processors that can be used for any workload and a large number of standby processors that can be used at no charge in the event of a disaster.

On iSeries, partitioning is enabled via the system kernel called a hypervisor, which allows and can control multiple OS types and instances on the same physical hardware. Each LPAR looks like a complete system.

Values derived from LPAR include:

- ▶ Multiple OS instances on the same hardware
- ▶ Multiple OS levels on the same hardware
- ▶ Workload separation
- ▶ Added availability characteristics

The iSeries i5 servers and the i5/OS V5R3, available since 2004, are based on the new POWER5 technology processor and the POWER-Hypervisor and they offer a new flexibility using the virtualization facilities allowed by these technologies. In particular:

- ▶ They support multiple operating systems, that facilitate server consolidation; they can run i5/OS, Linux, AIX 5L, and Windows on a single server
- ▶ They are designed to pool resources and optimize their use across up to 254 partitions running concurrently multiple application environment and operating systems

### **IBM @server xSeries**

This brand uses a different approach to on demand capabilities. On the high end (x445 and x455), the modular design of the systems allow for easy upgrades to the power of the system. These models can go from one to four Central Electronic Complexes (CECs); these complexes can operate standalone to allow four different systems, and when power requirements increase, it is then possible to interconnect up to four of these CECs to better handle the workload. When the peak requirements have passed, it is then possible to revert back to independent systems.

### **IBM @server BladeCenter**

IBM Standby Capacity on Demand for BladeCenter offers unprecedented flexibility and speedy deployment. The system ships with seven blades standard and seven standby blades to deliver active and reserve computing power. There is no need to wait for new blades to come from the factory and no need to disrupt current operations. Blades can be activated through the management console as needed over a six month period, based on business needs and business growth. Payment is tied to the capacity usage with no price premium.

## 16.5.2 Self-configuring IBM TotalStorage

IBM offers capacity on demand solutions that are designed to meet the changing storage needs of rapidly growing e-businesses. Standby CoD is designed to provide the ability to tap into additional ESS storage and is particularly attractive in environments with rapid or unpredictable growth.

IBM Standby Capacity On Demand for ESS (SCoD) provides “standby” storage for the ESS and allows access to extra storage capacity whenever the need arises. With Standby CoD, IBM will install up to six Standby CoD Disk Eight-Packs in an ESS for a nominal charge. At any time, they can be logically configured for use, a non-disruptive activity that does not require intervention from IBM.

## 16.5.3 Hypervisors and workload managers

The hardware products/features just described are important, but by themselves do not provide a complete solution. This hardware needs to be automatically provisioned and the operating systems and applications need to be able to take advantage of increased capacity transparently.

A *hypervisor* program is a very specialized program, entrusted to carefully perform operations such as virtual memory management, access I/O address space, DMA addressing, and resource isolation. The hypervisor program is stored in a system flash module in the server hardware. During system initialization, the hypervisor is loaded into the first physical address region of system memory. The hypervisor program is trusted to create partition environments, and is the only program that can directly access special processor registers and translation table entries.

Partition programs have no way to access the hypervisor's instructions or data, other than through controlled hypervisor service calls that are part of the processor architecture. These protections allow the hypervisor to perform its duties in a simple and rigorous manner, resulting in the confinement of each operating system to a very tight, inescapable box.

The hypervisor function can control how much resource is provided to every partition and dynamically make changes to support the current requirements.



Such hypervisors exist for the IBM @server zSeries, iSeries, and pSeries. For the xSeries brand, the same function can be fulfilled by third party software such as VMware, as shown in Figure 16-10. VMware allows for multiple virtual machines to be defined and installed in one physical environment.




Figure 16-10 VMware ESX architecture

Once the operating system is limited to its partition, enforced either by a hardware hypervisor or software such as VMware, multiple applications can be started in their own virtual systems. To ensure efficient use of the available physical resources and dynamically adjust them to meet the business needs represented by the applications, requires a workload manager component. The workload manager assigns the optimal partition resources to each application so it can reach its goal based on the business policy. The capabilities of the workload manager may vary with the operating systems involved.

Grid technologies, such as those provided in the IBM Grid Toolbox, can take advantage of job scheduling to spread complex tasks across heterogeneous resources.

Figure 16-11 shows an example of an advanced implementation of these concepts on an IBM *zSeries* system. It shows how the physical resources are first virtualized by the hypervisor, then split into four different partitions. In partition 1, we then have the workload manager which decides how much resource, such as CPU and storage, should be given to each application to meet its goal. The workload manager will constantly monitor the system and adapt processing and resources to meet the goals. On partition 4, we see additional virtualization done by z/VM, which is able to create independent operating systems (Linux in our scenario) and share its resources between each of the Linux images.

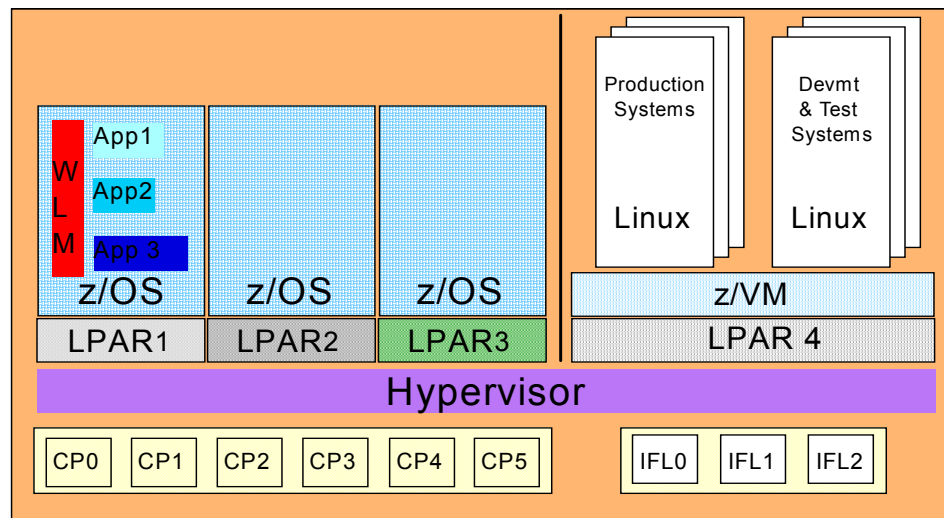


Figure 16-11 Use of resources on a zSeries

Linux with z/VM provides specific features to virtualize the network:

- ▶ A virtual LAN (VLAN) allows a physical network to be divided administratively into separate logical networks; in effect, these logical networks operate as if they are physically independent of each other.
- ▶ The Guest LAN, introduced with z/VM 4.2, allows the creation of multiple virtual LAN segments within a z/VM environment. With z/VM 4.3, the Guest LAN can be defined to use either HiperSockets devices or QDIO devices. With z/VM 4.4 z/VM Guest LAN support VLAN functions, that means Guest LAN can pass VLAN-tagged frames between guests that are VLAN-aware.
- ▶ z/VM Virtual Switch (VSWITCH, announced in z/VM V4R4, is an extension of the Guest LAN and is designed to improve the interaction between guests running under z/VM and the physical network connected to the zSeries. It provides a way to link an external network to guests under z/VM via an OSA Express card, without the need for a routing function. Guests attached to the

VSWITCH appear to be attached to the LAN that the OSA EXpress is attached to.

VSWITCH can be used to create highly-available connectivity for Linux guests under z/VM. z/VM guests can be accessed by configuring the redundancy features of VSWITCH combined with LAN-based high availability. Multiple OSA Express adapters can be defined for hardware redundancy, with multiple TCPIP controller service machines. As a result, as long as the LAN switch is configured appropriately, one can ensure that the z/VM guests stay linked to the external network when failures occur.

Figure 16-12 illustrates the ability of the IBM @server iSeries to optimize the IT environment through the consolidation of several infrastructure servers on a single machine. In addition to OS/400, the iSeries can host Linux and Microsoft Windows, as well as AIX in the future. (The possibility of AIX being supported in the future is a statement of direction. Such plans are subject to change by IBM without notice.)

The iSeries can provide both virtual and direct resources to this environment. With virtual I/O, the I/O resources (disk, tape, and CD-ROM) are owned by an OS/400 partition, which can share them with the non-OS/400 partitions. In addition, Virtual Ethernet can provide a secure, high-performance bus interconnect between partitions without an external LAN. These connections can provide 1 Gb performance and do not require LAN adapters, switches, or physical networks.




Figure 16-12 iSeries IT optimization

## 16.5.4 IBM Tivoli Provisioning Manager

Now that we have seen how an infrastructure might self-configure to adapt to changing conditions, we need to go a step further to where we have tools to monitor the environment and take the appropriate actions to adapt to those changes.

The software component which is going to start the action to modify the resources for each environment is the IBM Tivoli Provisioning Manager. It provisions and configures servers, operating systems, middleware, applications, and network devices acting as routers, switches, firewalls, and load balancers.

IBM Tivoli Provisioning Manager, through workflows, automates the manual provisioning and deployment process. It uses pre-built “industry best practice” workflows to provide control and configuration of major vendors' products. User-customized workflows can implement a company's data center “best practices” and procedures. These procedures can then be automated and executed in a consistent error-free manner.

## 16.5.5 IBM Tivoli Intelligent ThinkDynamic Orchestrator

The companion for IBM Tivoli Provisioning Manager is the IBM Tivoli Intelligent ThinkDynamic Orchestrator product.

Ever expanding, distributed islands of processing power need to be tapped to improve resiliency and responsiveness to a demand curve that can spike with only as much warning as an unexpected news release. Enterprises looking to reduce costs often cannot afford to maintain server capacity for just-in-case situations.

IBM Tivoli Intelligent ThinkDynamic Orchestrator helps them deal with these challenges. It helps reduce costs and improve server utilization. It helps boost server-to-administrator ratios by automating the steps to provision, configure, and deploy a solution into production. This automated process supports servers, operating systems, middleware, applications, and network devices acting as firewalls, routers, switches, and load balancers. By utilizing existing hardware, software, and network devices without rewiring, implementation times can be minimized, resulting in a faster return on investment.

Under-utilized assets can be put to work to help improve service levels with dynamic, on demand provisioning. Service levels can be constantly monitored; resource requirements anticipated for peak workloads and then automatically deployed. IBM Tivoli Intelligent ThinkDynamic Orchestrator provides a powerful system that:

- ▶ Protects existing investments, lowers implementation costs, and creates a fast return on investment by utilizing existing hardware, software, and network devices without rewiring or changing the network architecture.
- ▶ Helps reduce costs, improve server utilization, and boost server-to-administrator ratios by automating all the steps necessary to provision, configure, and deploy a complete solution into productive use
- ▶ Increases IT resource utilization tied directly to business results. Orchestration allows companies to manipulate their IT environment in real time — according to defined business policies — to achieve desired business goals. Orchestration does this by sensing an increase in the demand for resources and automatically taking action to re-allocate resources accordingly, and by provisioning resources throughout the entire system.
- ▶ Helps reduce security exposures by tracking and applying security patches to distributed network servers.
- ▶ Delivers a comprehensive business policy-based solution to automate provisioning, configuration, and deployment processes to use IT resources more efficiently.
- ▶ Anticipates, plans, and dynamically provides server capacity to meet peak business demands on demand, or in the case of system failures or outages

## 16.6 Linkages

The “packaged intelligence” gained from customer engagements is transformed into IBM infrastructure management solutions in the form of policy-based workflows. These workflows help the underlying IBM Tivoli Intelligent ThinkDynamic Orchestrator product make smart decisions on how to orchestrate the IT environment. The sequences of automated best practices articulate the administrative and management steps designed to integrate the business processes and priorities of an organization, and help the IT environment reach the desired state of operations.

As one example of how the various products that may make up a provisioning solution work together, let’s briefly look at the IBM Web Infrastructure Orchestration product.

To increase the efficiency and speed of provisioning from bare metal up to the running application, IBM Web Infrastructure Orchestration leverages IBM Director functionality in the automated best practices that control the IBM @server BladeCenter hardware, unleashing the power of BladeCenter and providing comprehensive remote management from a single graphical console.

The IBM Web Infrastructure Orchestration solution also helps protect a customer's critical data through the use of IBM Tivoli Storage Manager and IBM TotalStorage Solutions. This component manages and protects an organization's data from hardware failures and other errors by storing backup and archive copies of critical data with advanced replication services, delivering business continuity and facilitating effective disaster recovery.

Figure 16-13 shows the linkage and relationships between the orchestration products, which have the following functions:

- ▶ **IBM Intelligent ThinkDynamic Orchestrator:** Senses the environment and determines what changes need to be made based on a "Desired State Workflow".
- ▶ **IBM Tivoli Provisioning Manager:** Selects and executes the elemental provisioner to make the appropriate changes.
- ▶ **Drivers:** Drivers are supplied with Tivoli Provisioning Manager, and are used to execute the workflow actions and provide a way to interface with the resources.
- ▶ **IBM Director:**
  - Provides the interface between the workflows and the IBM hardware resources.
  - Removes the need to duplicate the interface and capabilities within the provisioning manager drivers.
  - Has the detailed information required to configure and manage the resource itself.
  - Understands the BladeCenter architecture and technology.
- ▶ **IBM Server Allocation for WebSphere Application Server:** Provides orchestration of workload and capacity optimization for WebSphere environments.
- ▶ **IBM's Tivoli Storage Manager and IBM TotalStorage Solutions:** Manages and protects an organization's data from hardware failures and other errors by storing backup and archive copies of critical data with advanced replication services.

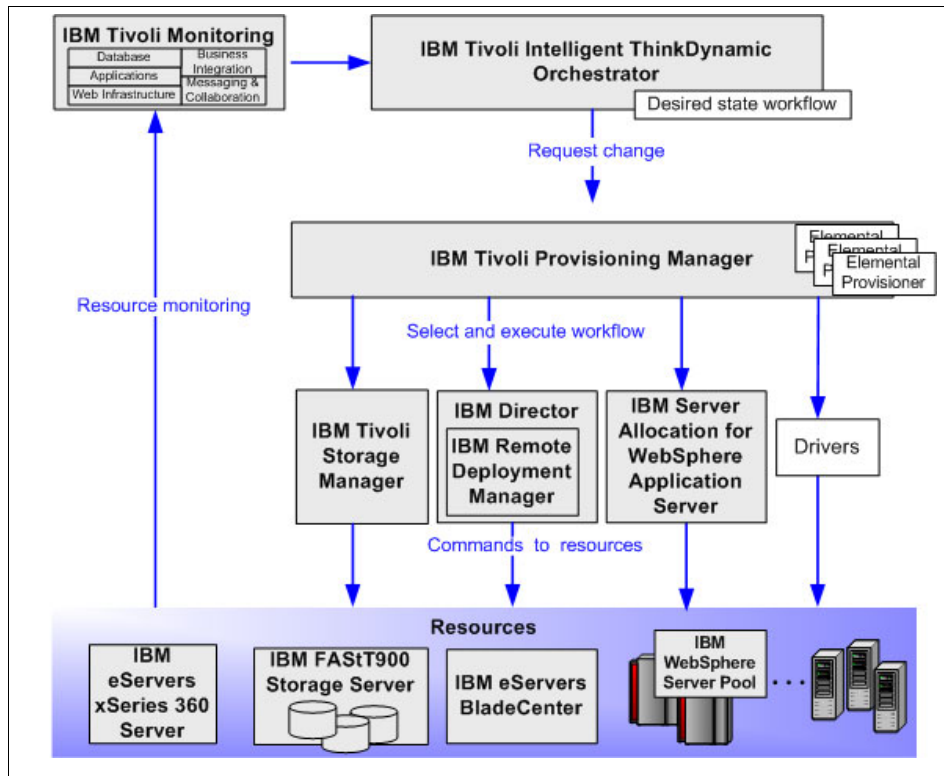


Figure 16-13 Linkage of provisioning products

## 16.7 Glimpse of the future

As technologies and products evolve, some of the items we might expect to see include having the same flexibility on all hardware platforms, regarding processor (including suppressors), memory, and I/O allocation, and the number of independent operating systems hosted.

The workload manager that today runs within the operating system may find its way to the hypervisor to be able to balance resource across partitions, and eventually across different systems and architectures.

The software components of provisioning solutions should also adapt to these changes, as they become able to interface with the global workload manager and improve its capability to define more complex enterprise policies.

Another important aspect of the On Demand Operating Environment related to provisioning is billing and metering capabilities. The ability to automatically measure the resources consumed by task, user, or department is critical to ease the billing process.

Finally, capabilities to tie in resources from on demand pools, as well as simpler configuration (self-managing autonomic configuration) for resource definition will be added. Systems will begin to discover their own “provisionable” resources, and make those available through self-configuring mechanisms.

## 16.8 Summary

In this chapter we have described a business issue centered around automated provisioning of systems to support applications, and more importantly, business requirements. We have described some of the products available today that could apply to businesses wanting to address similar issues. We have also described a scenario where some of these products are used together to provide a solution.

Taking a focused approach to meeting today's business needs by adopting products and solutions that fit into the On Demand Operating Environment framework allows customers to evolve their environment in a stepwise fashion. They can adopt individual components of the overall framework as needed to address business needs and feel comfortable in the knowledge that the On Demand Operating Environment will be an enabler for them becoming more responsive to the needs of their business and their customers.





## How to balance workloads in the network

This chapter describes how to use a Content Switching Module and the IBM Enterprise Workload Manager for efficient load balancing. It briefly explains the functionality of each component and details the configuration needed for both.

## 17.1 Introduction

System and workload management frameworks have proven to be valuable utilities used to monitor and manage enterprise applications and the systems they run on. While these systems may have a great deal of knowledge about these applications, they generally have little ability to control the rate at which they receive work.

In contrast, load balancers can control this rate, but typically have little information about the application's ability to successfully handle the request. The network balancing described in this chapter is a cooperation of IBM's workload management product, EWLM (Enterprise Workload Manager), and CISCO's load balancing product, the CSM (Content Switching Module), to provide an effective load balancing solution based on application performance and the ability of the applications to achieve business level goals.

To facilitate interaction between server load balancers such as the CSM and a workload managing entities such as EWLM, the Server/Application State Protocol (SASP) was developed. EWLM monitors the health and load of the servers and their applications in a configured cluster and makes decisions as to which servers or applications are best suited to handle client requests successfully. As part of this monitoring process, EWLM assigns a relative weight to each server in a cluster and sends these weights to the load balancer. The load balancer can then use these weights to balance client traffic to the most appropriate server.

The first Cisco load balancer to support SASP is the Content Switching Module (CSM), and the first IBM products to communicate with the CSM using SASP are Enterprise Workload Manager (EWLM) and the z/OS Load Balancing Advisor<sup>1</sup>. This document will only describe the interaction between EWLM and the CSM, please consult the appropriate z/OS documentation for similar information about the interaction between the z/OS Load Balancing Advisor and the CSM. The CSM gets server state information via a single, dedicated connection to the EWLM. SASP not only allows EWLM to provide relative weights for servers, but also allows the CSM or EWLM to remove a server from service. This may happen when the server is either configured out of service or when a failure is detected for that server.

---

<sup>1</sup> A new z/OS Load Balancing Advisor, a z/OS component available with z/OS V1.7 makes sysplex information available to network-based load balancers so that they can make better load balancing decisions. This helps protect busy target servers in a sysplex from being overloaded with new requests when they are already in danger of failing to meet their WLM service class goals or lack displaceable capacity. By helping meet WLM goals in a sysplex while allowing you to take advantage of network-based load balancers, z/OS Load Balancing Advisor helps maximize availability and optimize performance.

Figure 17-1 shows a diagram of the interaction of the CSM and EWLM.

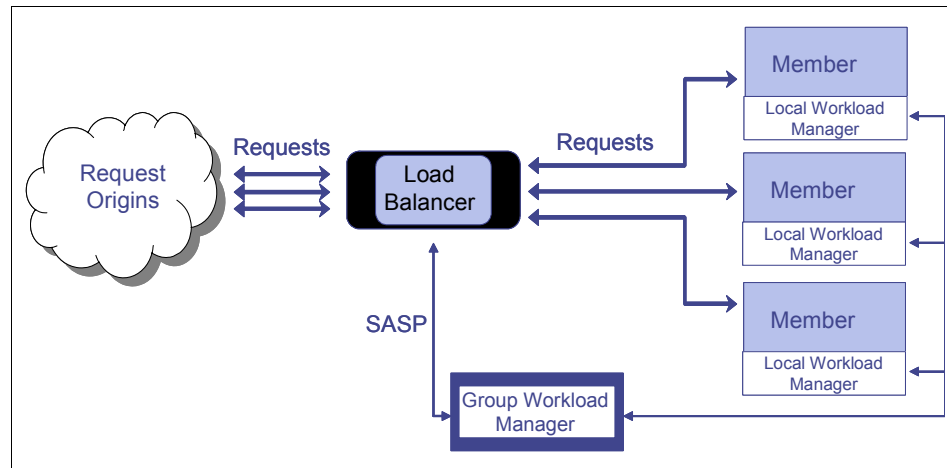


Figure 17-1 CSM and EWLM interaction

### 17.1.1 Overview of the CISCO CSM

The CSM distributes new connections to the server farm according to an administrator selected algorithm. If round robin or least connections are chosen as the load balancing algorithm, a weight will also be associated with each server in the server farm to allow the administrator or third party managers (such as EWLM) to provide the load balancer with distribution characteristics. When balancing according to one of these weighted algorithms, the CSM will balance client traffic to the servers so that the ratio of the number of client connections serviced by each server matches the ratio of the weights.

For example, if three servers, A, B, and C have weights 2, 3, and 1 respectively, the order in which traffic is distributed to the servers would be: A, B, C, A, B, B.

Note that during initial server farm creation, the CSM assigned a static weight to the server. By default, this weight is 8, but can be set to a different value by the administrator.

For a server farm in which EWLM has been associated, the weights which the CSM receives from EWLM will override any static weights. However, in the event that the EWLM becomes unreachable (due to either network failure or system failure), the CSM will use the weights that were configured when the server was first created in order to perform the load balancing scheme.

## 17.1.2 Overview of EWLM's load balancing recommendations

EWLM's Domain Manager has a bird's-eye view of the application topology to which the load balancer distributes traffic. The Domain Manager uses the application-level and system-level statistics it gets from the applications of the transaction topology to compose weights for the server farm describing the distribution that the load balancer should use for the next interval. These weights are formed to increase the likelihood that this work will be completed successfully while meeting the performance goals set by the administrator.

EWLM computes new weights for the server farm every 30 seconds. Some of the application and system factors used in the computation include:

- ▶ Remaining capacity
- ▶ Current application failure rates
- ▶ History of the application to meet its goals
- ▶ Importance level of the work currently running
- ▶ The amount of time the application spends blocked

Note that all of EWLM's statistics come from its Managed Server component that must be installed on the machines of the server farm. If this component is not running, EWLM will think that the machine is down and return a corresponding weight of 0. The CSM will treat this weight as 0 as an indication that the server is no longer in service and will remove it from load balancing consideration (see "Configuring EWLM for load balancing" on page 261).

## 17.1.3 The Server/Application State Protocol (SASP)

EWLM sends weights to the load balancer using the Server/Application State Protocol (SASP). SASP is a binary protocol comprised of several request-reply interactions. Control of these protocol interactions remains at the load balancer unless the load balancer explicitly releases its control. The message types of SASP are briefly described below:

- ▶ **Register Request / Reply:** Used by the load balancer or application member to register an application with EWLM.
- ▶ **DeRegistration Request / Reply:** Used by the load balancer or application member to de register an application with EWLM.
- ▶ **Set Member State Request / Reply:** Used by the load balancer or application member to quiesce or reactivate members of the server farm.
- ▶ **Set Load Balancer State Request / Reply:** Used by the load balancer to configure the SASP interaction and allow application members to use registration, de registration, and set member state messages.
- ▶ **Get Weights Request / Reply:** Used by the load balancer to get the current weights for the members of the server farm.

- **Send Weights Message:** This is the only message sent from EWLM without first receiving a corresponding request. If the load balancer configures EWLM to do so, EWLM will send this message containing the newest weights when it deems necessary.

## 17.2 Configuring the components

Configuring EWLM to work with the CSM and vice versa is normally done after setting up the EWLM management domain and the CSM load balancing domain. A typical EWLM management domain setup procedure will look like this:

1. Install and configure EWLM Domain Manager.
2. Configure EWLM Control Center.
3. Install and configure EWLM Managed Server on each managed node.
4. Enable ARM instrumentation in supported middleware.
5. Create domain policy to set business goal for each managed application.

Please refer to the EWLM section of the IBM *@server* Software InfoCenter for complete instructions on setting up EWLM.

[http://publib.boulder.ibm.com/infocenter/eserver/v1r1/en\\_US/index.htm?info/icmain.htm](http://publib.boulder.ibm.com/infocenter/eserver/v1r1/en_US/index.htm?info/icmain.htm)

A typical CSM load balancing domain setup procedure will look like this:

1. Perform basic setup of the Catalyst switch.
2. Connect the servers to the appropriate ports.
3. Create VLANs depending on applications needs.
4. Configure routing on the switch to allow communication to and from the servers.
5. Logically group servers together as server farms based on the applications they are serving.
6. Define virtual IP addresses for incoming traffic to reach these server farms.
7. Choose a load balancing algorithm for each server farm and optionally assign static weight to real servers.
8. Optionally define probes and sticky group policies.

Please refer to the Cisco CSM user guide for complete instructions on setting up CSM and load balancing.

We will now focus on the changes to both configurations to facilitate EWLM Load Balancing with the CSM.

Please note that Cisco CSM literature may use the term GWM (Group Workload Manager) to refer to any software that can give dynamic weights to members of a server farm using the SASP protocol. In this case, EWLM is a type of Group Workload Manager.

## 17.2.1 Configuring the Catalyst 6509 and the CSM for EWLM

This section describes what must be done to an existing valid CSM configuration to enable EWLM to provide load balancing weights for the server farm.

1. First, make sure you have the right CSM firmware release, 4.1.2 or higher, and a supported IOS release.
2. Configure an agent for the SASP EWLM connection.
3. Associate that agent with the appropriate server farm(s).

### Configuring an agent for the SASP/EWLM connection

A special DFP (LocalDirector's Dynamic Feedback protocol) agent must be configured, signified by a special bind identification number. If this bind id falls into the range of bind ids allocated for SASP communication, then the CSM knows that the connection is to use SASP for communication.

Example:

```
Router(config-slb-dfp)# agent 64.100.235.159 3860 65520
```

Where the syntax is:

```
Router(config-slb-dfp)# agent <ip address> <port> <bind id>
```

The CSM identifies the bind id as a GWM bind id through the use of environment variables. The environment variables SASP\_FIRST\_BIND\_ID and SASP\_GWM\_BIND\_ID\_MAX identify the first GWM bind id and the maximum number of bind ids, respectively. So, if a DFP agent is created with a bind id between SASP\_FIRST\_BIND\_ID and SASP\_FIRST\_BIND\_ID + SASP\_GWM\_BIND\_ID\_MAX, the agent will use SASP as the communication protocol to the GWM.

The CSM configures the SASP parameters to have EWLM automatically send server weights to the CSM whenever a change of weights occurs. However to ensure that communication with EWLM has not been lost, if weights are not received after a keepalive interval, the CSM will send a load balancer state message to EWLM. If the overall system is performing correctly, the CSM will receive a load balancer state response message from EWLM. This keepalive interval defaults to 180 seconds, but can be configured when the connection is created.

Example:

```
Router(config-slb-dfp)# agent 64.100.235.159 3860 65520 120
```

Where syntax is:

```
Router(config-slb-dfp)# agent <ip address> <port> <bind id> <keepalive>
```

### Associating the SASP agent with a server farm

Once the EWLM connection is initialized, the server farm can then be associated with the SASP/ EWLM Agent. Using the bind id assigned to the SASP/ EWLM agent, the server farm becomes bound to EWLM.

At this point, the CSM will register the servers in the server farm with EWLM and query EWLM for weights representing the server state. The bind ID represents a connection with the EWLM Domain Manager, so more than one server farm can use the same bind ID if each server farm is managed by the same Domain Manager.

Example:

```
Router(config-slb-sfarm)# bindid 65520
```

At this point, the health of the real servers in the server farm will be adjusted according to the feedback received from EWLM.

**Note:** Make sure the associated virtual server has an IP protocol and port set in order for EWLM to map the incoming requests to a particular PID on each real server, allowing application-level statistics to be made available to calculate better weights. If the protocol or port is not set, or the target application is not ARM instrumented, then weight calculation will still work but will be based on system-wide statistics instead of application related statistics.

## 17.2.2 Configuring EWLM for load balancing

Once you have a working EWLM management domain, to enable dynamic weight generation for load balancing, you only need to specify a listening port for SASP messages, as follows:

```
changeDM [.bat |.sh] workingDir -lbp xxxx
```

Where xxxx is a valid TCP/IP port number, or “Off” to disable weight generation

- At the startup of EWLM Domain Manager, if a valid load balancing port is specified, the weight manager will be activated and it will be waiting for connections from load balancers.

- ▶ Once connected and valid groups are registered, the Domain Manager will send weight updates to CSM whenever there is a change.
- ▶ If the Domain Manager detects that a Managed Server is no longer online or the target application is down, it will also send CSM a weight of 0 indicating that the real server is offline and no traffic should be forwarded to it.

### 17.2.3 Failover considerations

This section covers some considerations regarding failover.

#### CSM failure

Two CSMs configured in an active/standby configuration can guarantee synchronized server state information by connecting to the same EWLM to receive identical dynamic weight updates. However, in order for this configuration to work, the two CSMs must be distinguishable to the EWLM. Thus, they must hold different values in their SASP\_CSM\_UNIQUE\_ID fields. For example, one CSM can have the SASP\_CSM\_UNIQUE\_ID value of “CSM-1” while the other CSM would require a different value, such as “CSM-2”.

#### EWLM failure

On the CSM, only a single EWLM is expected to be associated with a given server farm. In other words, no backup EWLM is configured. Consequently, if the CSM loses connectivity to the EWLM, it is no longer able to obtain dynamic weights for servers. Thus, as described in the earlier sections, the CSM will use the statically configured weights in order to perform the appropriate weighted load balancing algorithm.

It is important to note that the CSM will continually attempt to contact the EWLM in the background. If the EWLM communication is re-established, it is able to immediately use the dynamically learned weights. This attempt to re-establish a broken connection channel happens once every 20 seconds.

## 17.3 Network balancing example

This section provides an example topology and configuration used to run a test load balancing scenario. In order to load balance the incoming traffic to Web Servers, there must be at least one CSM, and two or more Web Servers. The Web servers would then be connected to one or more WebSphere Application Servers and on the back-end, most likely, a database; either shared or distributed.



### 17.3.1 Network and application topology

This scenario consists of a CSM balancing over four transaction paths each containing a Web server (IBM HTTP Server), an application server (WebSphere Application Server), and a database (IBM DB2). Two of the paths consist of AIX machines, the third path uses Windows machines, and the fourth path uses Solaris systems. The EWLM Domain Manager is running on a dedicated Linux machine. The load balancer is a CSM module in the 3rd slot of a CISCO Catalyst 6509 Switch. Note that the configurations included in this document assume that the reader has installed the Virtualization Engine fix pack 1.1.020 (Figure 17-2).

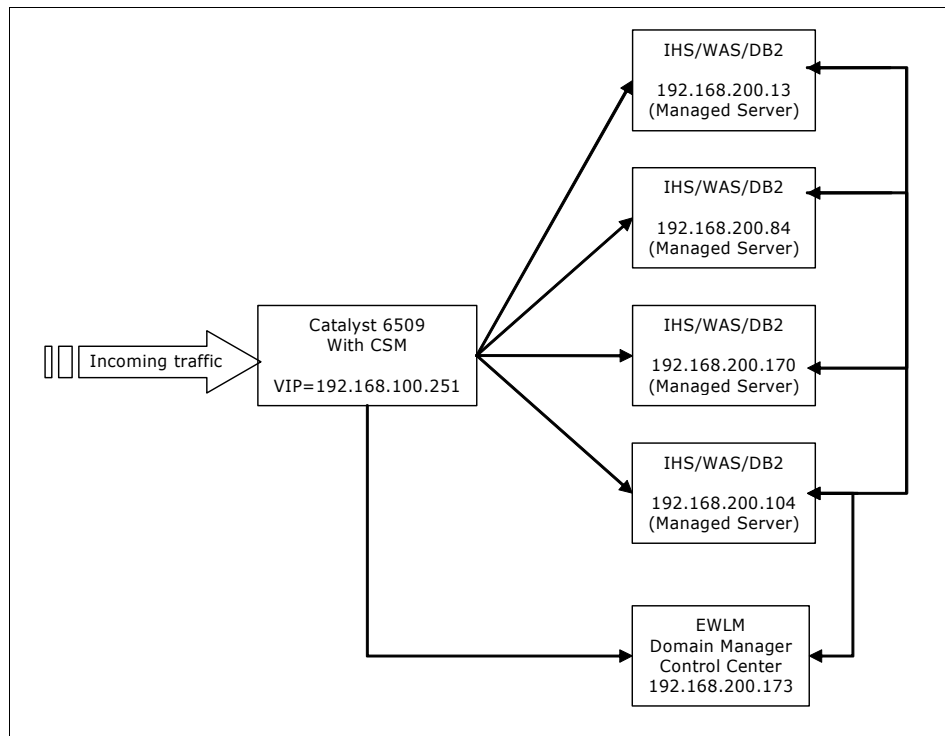


Figure 17-2 Typical network and application topology

### 17.3.2 Load balancing configurations

There are two configurations we discuss further in the section: EWLM and CSM.

## EWLM configurations

There are several EWLM control center panels that can verify that the EWLM management domain is working properly:

1. Managed Servers are connected to the Domain Manager. In the first EWLM screen shot in Figure 17-3 we can see which Managed Servers are currently connected and being used.

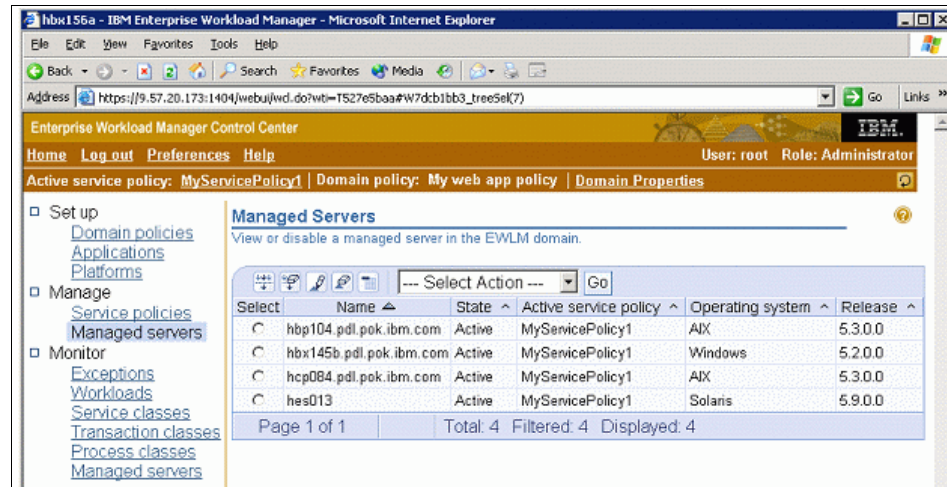


Figure 17-3 EWLM Control Center

In Figure 17-4, we can see the environment setup of this Case Study. You may notice in this configuration example that each hop in the transaction flow (IHS → WebSphere → DB2) is on the same machine (same IP address). This made the example easier, but certainly does not have to be the case.

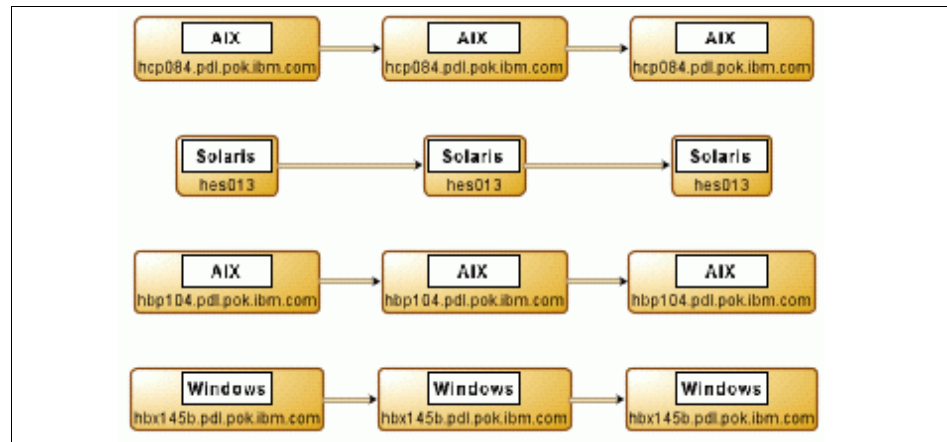


Figure 17-4 System topology view

2. All the supported middleware is ARM instrumented.
3. EWLM Control Center is showing transactions statistics.

**Note:** Items #2 and #3 above are not an absolute requirement, but would greatly improve the weight calculation algorithm if available.

These are the steps to configure EWLM Domain Manager to listen and respond to SASP messages:

1. Stop the Domain Manager.
2. Check the Domain Manager configuration:  
`./displayDM.sh /us/EWLMDM`
3. Add the load balancing port in the configuration table:  
`./changeDM.sh /us/EWLMDM -lbp 4447`
4. If your Domain Manager has 2 IP addresses, and you need to use both IP addresses for Managed Servers and/or load balancing, make sure your “-ma IP” parameter is 0.0.0.0 instead of a specific IP address, because “-mp xxxx” and “-lbp xxxx” ports will both use “-ma IP” parameter as IP address to bind to. To change that, use this command:

`./changeDM.sh /us/EWLMDM -ma 0.0.0.0`

You can combine the previous command and this one into one single command.

5. Start the Domain Manager again
6. Make sure Domain Manager is now listening on the load balancing port:

(command in windows) => `intestate -an`

Look for this:

TCP	0.0.0.0:4447	0.0.0.0:0	LISTENING
-----	--------------	-----------	-----------

## CSM configurations

All of these steps must be done as privilege 15, either by using “login” or “enable” commands at the console.

1. Verify the default SASP variables:

`sh mod csm 3 variable`

2. Change the value of a variable:

`configure terminal`  
`mod csm 3`

3. Create a server farm:

```
serverfarm testfarm
nat server
no nat client
predictor leastconns
bindid 65520
real 192.168.200.84
inservice
real 192.168.200.170
inservice
real 192.168.200.104
inservice
real 192.168.200.13
esvt6509(config-slb-real)#inservice
```

4. Create a virtual server:

```
vserver testvserver
virtual 192.168.100.251 tcp www
serverfarm testfarm
persistent rebalance
inservice
```

5. Create a DFP agent:

```
dfp
agent 192.168.200.173 4447 65520
end
```

6. Verify the serverfarm:

```
sh mod csm 3 serverfarms detail
```

7. Verify the real servers:

```
sh mod csm 3 reals
```

real	server farm	weight	state	conns/hits
192.168.200.84	TESTFARM	56	OPERATIONAL	0
192.168.200.170	TESTFARM	56	OPERATIONAL	0
192.168.200.104	TESTFARM	56	OPERATIONAL	0
192.168.200.13	TESTFARM	56	OPERATIONAL	0

8. Verify the virtual server:

```
sh mod csm 3 vservers detail
```

9. Verify the DFP agent:

```
sh mod csm 3 dfp detail
DFP Agent 192.168.200.173:4447 Connection state: Connected
Keepalive = 65520 Retry Count = 0 Interval = 180 (Default)
```

```
Security errors = 0
Last message received: 16:12:14 EST 06/22/04
Last reported Real weights for Protocol TCP, Port www
  Host 192.168.200.84   Bind ID 65520  Weight 56
  Host 192.168.200.170 Bind ID 65520  Weight 56
  Host 192.168.200.104 Bind ID 65520  Weight 56
  Host 192.168.200.13  Bind ID 65520  Weight 56
DFP manager listen port not configured.
No weights to report to managers.
```

## 17.4 Lessons learned

This section describes the lessons we learned while testing and running case studies of the CSM's usage of EWLM weights. It should be used as a reference for tips to provide more effective load balancing with EWLM and the CISCO CSM.

### 17.4.1 Best practices

1. It is recommended to start with a functional EWLM management domain and a functional CSM load balancing domain first, then enable the communication between EWLM Domain Manager and CSM.
2. In most of our tests, the weighted least connection algorithm yields better performance than the weighted round robin algorithm.
3. Beware of bottlenecks and single points of failure when using the EWLM Firewall Broker. A Firewall Broker acts as a proxy server, accepting connections from all the Managed Servers in the same IP subnet and channels them to the Domain Manager. As you can imagine, if the Firewall Broker machine or process is not available, all the Managed Servers relying on it will be disconnected from the Domain Manager.

Normally, the fact that the Managed Servers are disconnected from Domain Manager won't affect the functioning of any middleware; however, if EWLM load balancing is enabled at the CSM, the Domain Manager senses that a Managed Server is offline and sends a weight 0 to the CSM, stopping traffic from being forwarded to that server. If the Firewall Broker is suddenly unavailable, then the Domain Manager and CSM will believe that the entire server farm is unavailable.

4. The most reliable topology is to have the fewest hops possible between the Domain Manager and Managed Servers.

## 17.4.2 Special EWLM benefits to load balancing

EWLM load balancing weights help the CSM get better performance in typical load balancing scenarios. There are also several special scenarios where EWLM can provide exceptional benefit:

1. When the server farm contains servers of different capacities and hardware architectures in a single tier: EWLM can sense the application level capacity of these systems and adjust weights as they change.
2. When the server farm contains servers performing multiple types of work in addition to the work being load balanced: EWLM can be sensitive to the effect the other work has on the load balanced work. Furthermore, if the other work is properly instrumented, EWLM can try to manage the rate in which load balanced work arrives, to achieve business goals of all work on the system.
3. When the work running on the server farm machines has varying importance levels: EWLM is aware of the importance level of the work and will favor machines running less important work to preserve the resources on the systems which are currently running more important work.
4. When the work running on the server farm machines has specific goals to be met: EWLM is aware of these goals and will favor machines that have a history of achieving such goals.
5. When the applications receiving the load balanced work encounter failures (even downstream failures): EWLM is made aware of such application-level failures and will heavily favor machines not experiencing these failures. The applications experiencing such failures will continue to have minimal weights until EWLM is assured that they have recovered.

## 17.4.3 Troubleshooting

Here are some considerations for troubleshooting:

1. Both setting up a load balancing domain in the CSM and installing and configuring Enterprise Workload Manager in sophisticated enterprise environments could be quite complex. If load balancing problems arise in this environment, eliminating problem sources should begin with removing the DFP agent pointing to the EWLM Domain Manager.

Usually the Domain Manager would only change the weight of the real servers in CSM, but won't prevent traffic from flowing back and forth. The only time where it may prevent traffic from flowing is when the Domain Manager thinks that a Managed Server is offline or the ARM instrumented edge application is shutdown. In this case, the Domain Manager will send a weight of 0 to tell CSM not to send any more traffic to that application. Use the command "sh mod csm 3 reals" to see if you have any real servers in the "DFP\_THROTTLED" state (assuming that your CSM is installed in slot 3).

If you do have this problem, check in the EWLM Control Center to find out what's wrong. Then check in the Managed Server if everything is running properly, especially the Managed Server Java process, the middleware, and the network connectivity to the Domain Manager. Once you restore everything back to order, check in the EWLM Control Center to make sure that managed server is working properly again.

2. Logging may prove quite useful in troubleshooting SASP problems. Log messages are stored in the Catalyst's buffer, and the administrator may choose to use an external syslog daemon for more flexibility and storage options.

To display all the messages in the Catalyst's buffer, issue this command:

**show logging**

To set up logging to remote syslog, refer to the Catalyst's user guide and your operating system's manuals.

3. Another common problem seen in testing occurs when the DFP agent is not using the SASP protocol to communicate with EWLM Domain Manager. You will not see any logging messages at either end when this happens.

In the output of "sh mod csm 3 dfp detail," you'll see the DFP agent state as "Not connected", or "Trying to connect" (assuming that your CSM is installed in slot 3). To fix this, you need to check your configuration again, especially bind ID, and make sure it is in the range specified by SASP\_FIRST\_BIND\_ID and not higher than SASP\_GWM\_BIND\_ID\_MAX incrementally.

All the SASP error messages and return codes are logged in the logging buffer, or remote syslog if appropriately configured. Successful return codes are not logged.

## 17.5 Conclusion

Through advanced resource and statistics gathering techniques, EWLM can make intelligent weight assignments based on the relative load and availability of servers within a cluster. With the use of SASP, EWLM is then able to pass this relative state information on to the CSM, which in turn is able to balance client requests to the most appropriate server. The result is a well balanced, highly efficient distribution of traffic that ensures the best utilization of available resources as defined by the administrator.







## How to consolidate, simplify, and optimize the storage IT infrastructure

This chapter focuses on the need for heterogeneous storage infrastructure consolidation, simplification, and optimization, in order to map to the business requirements and changes, both better and faster. It focuses also on the need for centralized management tools especially designed for heterogeneous storage environments and that can provide advanced features to simplify and automate this management task.

A key component of the On Demand Operating Environment for this consolidation is virtualization, especially at storage level. One major aspect of the storage virtualization is the abstraction it procures between user and the application from the underlying physical resources that are being used and accessed. This virtualization enables implementation of advanced features such as *Policy-Based Data Placement*, *Policy-Based Life Cycle Management*, *Data Migration*, or *File Movement*, which can be performed both transparently regarding application server usage and across heterogeneous physical storage subsystems (IBM or non-IBM).

Another key component of the On Demand Operating Environment is the centralized storage management product set that can provide a global view of the full storage space and its usage by the users or applications.

This chapter identifies business needs as they relate to consolidating and simplifying the storage infrastructure in an On Demand Operating Environment. It identifies also the products available today from IBM that can be used in this solution to address these business needs. The products are mapped to the On Demand Operating Environment framework to show where they fit. A practical scenario is used to illustrate challenges that IT organizations face regarding storage resources.

## 18.1 Introduction

Storage consolidation and simplification in On Demand Operating Environment includes the capability to address the following business requirements:

- ▶ Improve the overall usage of the storage infrastructure by reducing the percentage of unused space.
- ▶ Increase resilience and eliminate potential Single Points Of Failure (SPOF's).
- ▶ Improve flexibility in the usage of the storage space by allowing cross platform non-disruptive data migration and cross-platform data compatibility.
- ▶ Also improve flexibility for future storage growth and expansion and create a storage infrastructure capable of growing in line with the business requirements.
- ▶ Reduce data duplication by providing a more effective data sharing between the heterogeneous server and applications.
- ▶ Reduce complexity of adding new elements to that storage infrastructure.
- ▶ Reduce the complexity of managing that infrastructure.
- ▶ Reduce the cost of managing the infrastructure.
- ▶ Provide a high degree of automation to simplify the infrastructure management and improve its resilience.
- ▶ Help in quickly deploying new applications and all of the required storage resources.
- ▶ Reduce time and cost to change or expand storage resources available for application to meet business changes.
- ▶ Reduce traffic on Local Area Network (LAN).

Different solutions or technology are available to allow the storage infrastructure to meet these business requirements; they will be detailed in the next sections.

## 18.2 General strategy

The components needed to implement this approach have to respect the general attributes of an On Demand Operating Environment. Specifically, they must be:

- ▶ **Flexible:** Solutions need to improve flexibility in the global storage space utilization and re-distribution regardless of the storage subsystem themselves, the application servers themselves, and the operating system.
- ▶ **Scalable:** Solutions need to be scalable to be able to follow and support the evolution in storage growth.
- ▶ **Resilient and recoverable:** The solution has to improve application or systems recovery time in case of failure or disaster. The components of the solution themselves have to be resilient and able to recover from any single failure on its elements.
- ▶ **Transparent for application:** The solution has to permit management operation to be, as much as possible, transparent for the application by minimizing or avoiding their down time.
- ▶ **Easy to manage:** The solution, even if it is based on technically complex elements or features, has to be easy to manage and configured with graphical and intuitive management interfaces.
- ▶ **Automatic:** The solution and its components has to provide the necessary features and interfaces or tools, allowing many basic management tasks to be performed automatically.
- ▶ **Economical:** Any solution today needs to be economical, and a storage solution is no different.
- ▶ **Centralized management:** The global storage infrastructure has to be manageable from centralized tools designed to support a heterogeneous storage environment.
- ▶ **Open standards-based:** The only way that a centralized solution can provide consistent management to a large number of platforms, operating systems, applications, and data stores, is by building on and using open standards.

## 18.3 Solution components

This approach focuses on both IBM Storage virtualization products and IBM TotalStorage Storage management software:

- ▶ The goal with Storage Virtualization products, SVC and SFS, is to improve and optimize storage resources utilization by adding a virtualization layer between the storage subsystems and the servers to address the limitations inherent to the standard SAN architecture and their static relationship between servers and storage subsystems.
- ▶ Centralized storage management software such as IBM TotalStorage Productivity Center (TPC) products have been designed to improve the management from a central point of a heterogeneous infrastructure and to provide different advanced features to optimize and automatize its utilization.

## 18.4 Scenario

The following sections describe the environment and the implementation of the solution.

### 18.4.1 Current environment

This existing customer infrastructure is composed of the following components:

- ▶ Open System servers and clients (UNIX, Linux and Windows).
- ▶ Data base applications running on a UNIX server connected to storage subsystems through a SAN fabric.
- ▶ File servers running on Linux or Windows servers connected to storage subsystems through a SAN.
- ▶ End clients PCs or servers connect to the database server or the file servers over an ethernet LAN.
- ▶ Heterogeneous storage equipment and subsystems (IBM or non-IBM) equipment. Some older subsystems with lower performances cannot be used for all applications and are reserved only for temporal storage space.
- ▶ A centralized backup/restore product which requires different agents installed on each server.
- ▶ Different workstations (PCs) dedicated to the management of the different parts of the storage infrastructure.

The existing customer storage infrastructure is described in Figure 18-1.

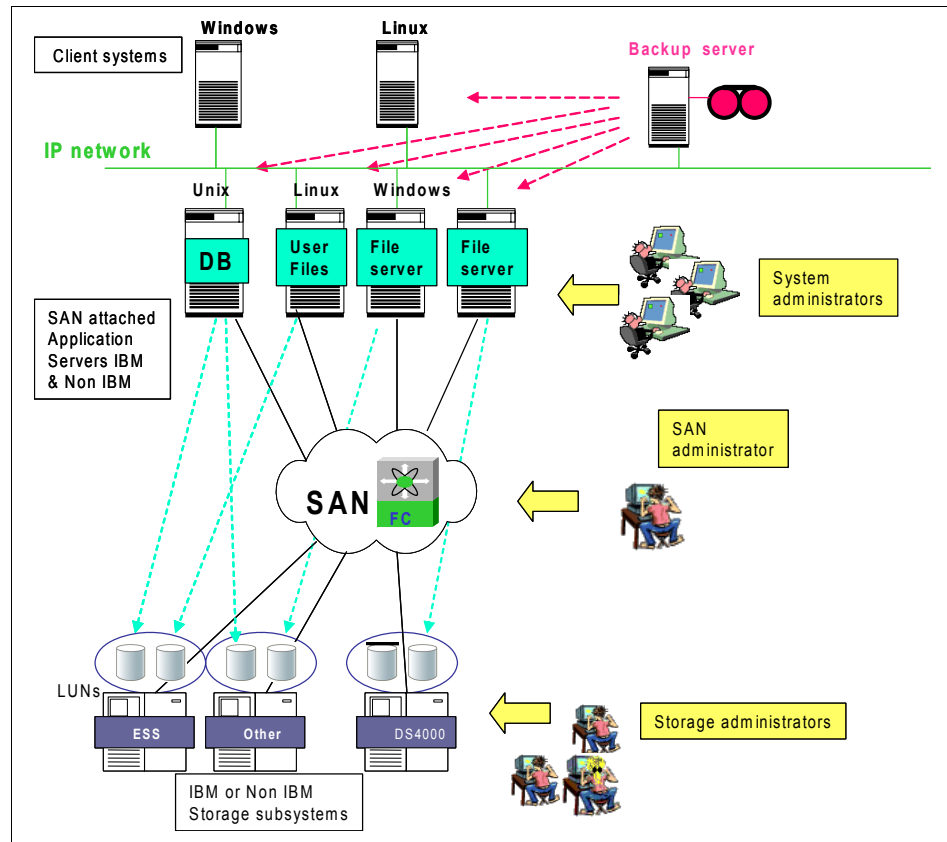


Figure 18-1 Current Customer Storage configuration

For the management of the storage infrastructure, several configuration tools and administrators skills are required:

- **Storage subsystems:** Each storage subsystem has its own management tool that can be accessible through a graphic user interface (GUI) or by a command line interface (CLI). Most of these tools required the installation of software on a management station (PC).

A GUI is more user friendly, in contrast, CLI requires a better knowledge from a specialized storage administrator, especially to create and maintain configuration or management scripts using this CLI interface. Each storage subsystems administrator needs to be available at all times, because their services are required each time a modification is requested, such as the creation of new LUNs and their attribution to an application server.

- ▶ **SAN Fabric:** SAN administration requires also the usage and the knowledge of different management tool. The SAN administrator is always requested each time a SAN modification is required; for example, in the zoning configuration.
- ▶ **Servers:** Several system administrators are also required, one per major operating system, to manage all the storage aspects. They have to closely coordinate with the other storage administrators if any change is required. They are also in charge of monitoring, with appropriate tools, the filling and usage of each logical disk to prevent and anticipate any outage (file system full, performance problems, and so on.)
- ▶ **Application:** Major applications such as databases or ERPs also have their administrator who must collaborate with system administrators and storage administrators for all questions and requirements regarding storage aspects.
- ▶ **Backup system:** An administrator is also required to manage the centralized backup system.

## 18.4.2 Business objectives

In response to the continuously growing need for storage space and to address existing problems or limitations, the company management team decided to achieve a major evolution in its storage architecture, first by making better use of existing storage space and avoid the necessity of adding new storage subsystems. Another major goal of this reorganization was to reduce the storage management effort and cost. Their requirements included the following points:

- ▶ Storage infrastructure has to be available 24 hours a day, 7 days a week.
- ▶ Existing storage space has to be consolidated to permit a better and fully usable by all applications, especially for the space residing on older and low performance subsystems.
- ▶ Flexibility in storage subsystems utilization and configuration has to be improved to always provide the optimum storage space utilization regarding business changes.
- ▶ Storage space utilization has to permit a better sharing data between heterogeneous servers by allowing true-sharing (only one copy of the same data) instead data replication.
- ▶ Time to meet business requirements has to be reduced, especially in case of peaks demands.
- ▶ Storage infrastructure management cost and time has to be reduced.
- ▶ Storage space utilization also has to be improved by the implementation of a Data Life Cycle Management strategy.

### 18.4.3 Technical objectives

As we can see, the company wants to get more from their existing IT infrastructures and do so within a reasonable cost. It needs an optimization of all the manual processes by a centralization of the management tools with a greater automatization of basic tasks. It also needs to optimize the storage space utilization in an new architecture that allows quick and non-disruptive adjustment for rapidly-changing business conditions, as well as keeping space for fluctuating demands of key business applications. Finally, a product allowing data Life Cycle Management has to be added to the storage infrastructure to permit older data to be automatically migrated in less expensive storage spaces.

The following list summarizes the technical objectives for the target storage infrastructure that has to be implemented to meet the required business objectives and solve several limitations of the existing storage architecture:

- ▶ Find a solution that could permit a full pooling of all the storage space located on the different heterogeneous storage subsystems as a global storage pool, and also optimize storage space usage and add flexibility in its configuration.
- ▶ Provide high performance data sharing between the heterogeneous servers.
- ▶ Introduce a product that allows the file placement control between the different available storage pool. This product has also to provide an automatic Life Cycle Management (LCM) feature.
- ▶ Introduce centralized storage management tools with integrated advanced feature that could permit a better automatization of the storage management and monitoring tasks.

### 18.4.4 Solution approach

IBM TotalStorage Virtualization products and IBM TotalStorage Productivity Center products, which are key components of the IBM TotalStorage software family, have been chosen to meet the business and technical objectives that were listed previously.

#### **IBM TotalStorage SAN Volume Controller**

IBM TotalStorage SAN Volume Controller (SVC) is the first IBM storage virtualization product. SAN Volume Controller, which is designed to help manage the complexity and costs of SAN-based storage. It allows a logical consolidation of the storage space at the disk block level. Based on virtualization technology, the SVC supports a virtualized pool of storage from the storage subsystems attached to a SAN. The main benefits of the SVC are a centralized control for volume management, the capacity to combine storage space (LUNs) of multiple heterogeneous storage controllers with a single view of the volumes, advanced data migration, and data copy services.

## IBM TotalStorage SAN File System

The IBM TotalStorage SAN File System (SFS) is another IBM storage virtualization product. SFS is intended to combine the benefits of file sharing across servers provided by the NAS architectures with the benefits of high performance data access provided by Storage Area Networks (SAN).

SFS is designed to provide network-based heterogeneous file system to support data sharing with policy-based file placement and file movement features (LCM) in an open environment. SFS also provides logical consolidation of the storage space such as SVC, but at the file level. Like SVC, it provides a central management point and implements advanced copy services.

Two other products of the IBM TotalStorage Productivity Center will be also used to manage the new IT storage infrastructure from a central point:

- ▶ **IBM TotalStorage Productivity Center for Fabric:** This product brings all the information about SAN topology and configuration into a single place and is able to create, by correlation between different sources of information, topology mapping of the SANs which are displayed graphically from a central management point. It can also provide other features such as automated device discovery, error detection fault isolation, SAN error predictor, zone control, real-time monitoring and alerts.
- ▶ **IBM TotalStorage Productivity Center for Data:** This product is designed to provide a comprehensive Storage Resource Management (SRM) solution by providing reports, monitoring and alerts, policy-based action, and file system capacity automation.

Figure 18-2 shows how these products are used in the new Storage IT architecture.



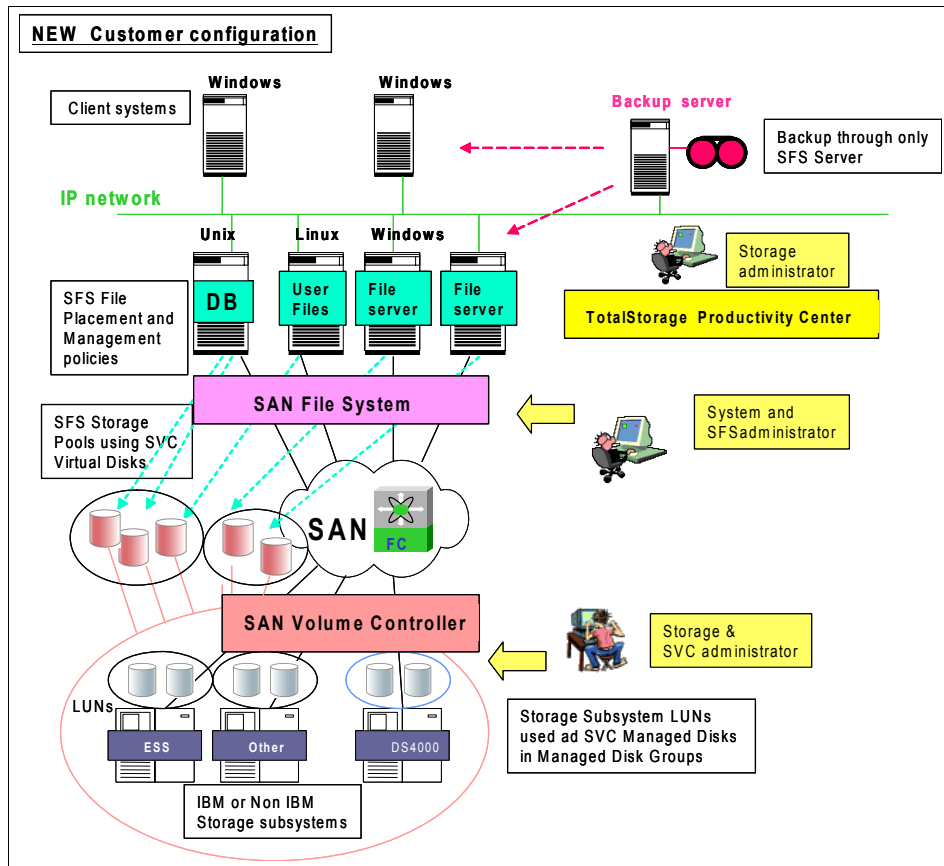


Figure 18-2 New IT Storage Architecture

## Configuration aspects

The following sections detail the implementation setup.

### IBM TotalStorage SAN Volume Controller (SVC) configuration

A logical consolidation of the storage space is performed by the SVC at the disk level.

To do that, all the storage subsystems LUNs must be no longer be directly mapped to the application servers, but now to the SVC nodes; for a 2-node SVC cluster, which is the minimum configuration, each LUN has to be mapped to 8 Fibre Channel ports (an SVC node has 4 Fibre Channel ports). These LUNs, which are named *Managed Disks* under SVC terminology, are then grouped in *Managed Disk Groups* by the administrator, who can then create the *Virtual Disks* from them, and map them to the application server.

Virtual disks are split (striped) between the different Managed disks of the Managed Disk Group with the possibility to choose from only one disk (sequential Virtual Disk) to all the disks. It is theoretically possible to put together any LUN in a Managed Disk Group. However, in practice, due to the striping mechanism, it is strongly recommended to use LUNs with the same characteristics and from the same storage subsystem in each Managed Disk Group.

For all these reasons, the SVC and storage subsystems target configuration, in this scenario, will be as follows:

- ▶ Each storage subsystem will be reconfigured to create bigger LUNs with, if possible, only one LUN per full storage array (physical block of disks). This has to be done with respect to the storage subsystem recommendations, especially for performance and load balancing reasons; on a dual controller subsystem such as the DS4000, it is recommended to create LUNs by pair, each one being handled by a different controller.
- ▶ Each new LUN will be configured with different RAID level RAID5, RAID1, or RAID0.
- ▶ All these LUNs, which are seen as Managed Disks by SVC when it discover them, will be grouped by the administrator in the following Managed Disks Groups:
  - a. ESS-RAID1-MDG: This group will include the ESS mirrored LUNs (RAID0) and will be used mainly to create Virtual Disks for the DB.
  - b. ESS-RAID5-MDG: This group will include the ESS Raid5 LUNs and will be used also to create Virtual Disks for the DB.
  - c. DS4000-RAID0-MDG: This group will include DS4000 Raid0 (striped) LUNs and will be used to create Virtual Disks for temporary files on this fast but not secured storage area.
  - d. EMC-RAID5-MDG: This group will include the EMC Raid5 LUNs and will be used to create Virtual Disks for different data files.
  - e. DS4000-RAID5-MDG: This group will include the DS4000 Raid5 LUNs and will be used to create Virtual Disks for all the other files (default storage pool).
- ▶ Virtual disks will be created in the different Managed Disk Groups to reflect the old storage configuration. In place of each storage subsystem LUN that was directly mapped to a server, we will create a Virtual Disk of same size and in a Managed Disk Group of same RAID level.

**Note:** To migrate from the existing storage configuration to the new virtualized SVC configuration, the existing LUNs could be integrated into the SVC configuration by using a specific import feature, called *Image Mode*, which avoids the lost of the data.

In “Setting up the environment” on page 286, a step by step SVC configuration is discussed. Example 18-1 details the SVC configuration.

#### *Example 18-1 SVC configuration*

---

##### **Managed Disks:**

```
IBM_2145:admin>svcinfolsmdisk -delim :
id:name:status:mode:mDisk_grp_id:mDisk_grp_name:capacity:ctrl_LUN_#:controller_
name
0:mDisk0:online:managed:0:ESS-RAID1-MDG:50.0GB:0000000000000000:ESS
1:mDisk0:online:managed:0:ESS-RAID1-MDG:50.0GB:0000000000000000:ESS
2:mDisk1:online:managed:1:ESS-RAID5-MDG:200.0GB:0000000000000001:ESS
3:mDisk2:online:managed:1:ESS-RAID5-MDG:500.0GB:0000000000000002:ESS
4:mDisk3:online:managed:2:EMC-RAID5-MDG:200.4GB:0000000000000003:EMC
5:mDisk4:online:managed:2:EMC-RAID5-MDG:200.4GB:0000000000000004:EMC
6:mDisk5:online:managed:3:DS4300-RAID5-MDG:400.2GB:0000000000000005:DS4300
7:mDisk8:online:managed:4:DS4300-RAID0-MDG:200.2GB:0000000000000006:DS4300
```

##### **Managed Disk Groups:**

```
id:name:status:mDisk_count:vDisk_count:capacity:extent_size:free_capacity
0:ESS-RAID1-MDG:online:2:2:100.0GB:32:40,0GB
1:ESS-RAID5-MDG:online:2:2:700.0GB:32:100,0GB
2:EMC-RAID5-MDG:online:2:1:400.0GB:32:300,0B
3:DS4300-RAID5-MDG:online:1:2:400.0GB:32:0,0GB
4:DS4300-RAID0-MDG:online:1:1:200.0GB:32:150,0GB
```

##### **Virtual Disk:**

```
IBM_2145:admin>svcinfolsvdisk -delim :
id:name:IO_group_id:IO_group_name:status:mDisk_grp_id:mDisk_grp_name:capacity:t
ype:FC_id:FC_name:RC_id:RC_name
0:DB-VDisk0:0:IO_grp0:online:0:ESS-RAID1-MDG:40.0GB:striped:::
1:DB-VDisk1:0:IO_grp0:online:0:ESS-RAID1-MDG:20.0GB:striped:::
2:DB-VDisk2:0:IO_grp0:online:1:ESS-RAID5-MDG:200.0GB:striped:::
3:DB-VDisk3:0:IO_grp0:online:1:ESS-RAID5-MDG:400.0GB:striped:::
4:Data-VDisk1:0:IO_grp0:online:2:EMC-RAID5-MDG:100.0GB:striped:::
5:Data-VDisk2:0:IO_grp0:online:2:3:DS4300-RAID5-MDG:50.0GB:striped:::
6:Default-VDisk0:0:IO_grp0:online:3:DS4300-RAID5-MDG:350.0GB:striped:::
7:Temp-VDisk0:0:IO_grp0:online:3:DS4300-RAID0-MDG:50.0GB:striped:::
```

---

As we see in this example, several virtual disks have been created in different Managed Disk groups using LUNs (Managed LUNs) from different storage subsystems and configured with different RAID levels. We will see later how they could be migrated to other Managed Disk groups, but for now, all of them will be used by SFS to create its Storage Pools such as those described next.

**IBM TotalStorage SAN File System (SFS) configuration**

Once the first logical consolidation will be done by SVC at the disk level, it will be now possible to perform the SAN File System configuration to get another logical consolidation, but now at a file level:

- ▶ First, we have to group the different SVC Virtual Disks into several SFS “Storage Pools”; the different files created into the SAN File system structure may be physically copied into these storage pools. As with SVC Managed Disk Groups, SFS Storage Pools will be defined by the administrator with the different available SVC Virtual Disks.
- ▶ The administrator will need also to define the *Automatic File Placement* policies; they determine, through their different rules, into which Storage Pool will be physically located each file created into the SFS Global name space. These rules, which have an SQL-like statement, accept different arguments such as file size, filename, file extension, creation date, etc. Some examples of possible rules are described in Example 18-2.
- ▶ The administrator has also to logically divide the SFS Global name space in several subsets named “Filesets” for the following reasons:
  - As each fileset is assigned to a specific SFS Meta Data server for its management, this produces a simple load balancing between the different Meta Data Servers (MDS) of the SFS cluster.
  - Although each Fileset will be seen by SFS client such as a standard directory; from a management point of view, some operations are only possible at the fileset level such as the SFS *Flashcopy* or the file sharing configuration.
  - Fileset names can be also used as supplementary argument in the File Placement policy rules; for example, we can define a rule for \*.mp3 files created in fileset Users.
- ▶ Finally, the SAN File Systems agents have to be installed on each server that will use the SFS

*Example 18-2 SFS configuration*

**Storage Pools:**

```
sfsccli> lspool
Name Type Size (MB) Used (MB) Used (%) Threshold (%) Volumes
=====
```

```

SYSTEM System 2032 240 11 80 1
DEFAULT_POOL User Default 354024 245226 69 80 1
DB-Pool1 User 612384 423678 69 95 2
DB-Pool2 User 612384 238677 39 95 2
Data-Pool1 User 158342 87652 55 80 2
Temp-Pool User 51245 12045 23 80 1

```

Filesets :

```

sfsccli> lsfileset
Name Fileset State Quota Type Quota (MB) Used (MB) Used (%) Threshold (%) Most
Recent Image Server
=====
DB Attached Soft 0 0 0 80 - mds0
Files Attached Soft 0 0 0 80 - mds1
Users Attached Soft 0 0 0 80 - mds1
Temp Attached Soft 0 0 0 80 - mds2
Programs Attached Soft 0 0 0 80 - mds2

```

### File Placement policies:

```

sfsccli> lspolicy
Name State Last Active Modified Description
=====
DEFAULT_POLICY inactive May 06, 2004 3:40:05 AM May 06, 2004 3:40:05 AM Default
policy set (assigns all files to default storage pool)
production_policy active - May 14, 2004 3:14:22 AM Production Policy

```

```

sfsccli> catpolicy production_policy
VERSION 1 /* Do not remove or delete this line! */
rule 'ProdRule1' SET STGPPOOL DB-pool1 FOR FILESET (DB) where NAME LIKE '%data%'
rule 'ProdRule2' SET STGPPOOL DB-pool1 FOR FILESET (DB) where NAME LIKE
'%index%'
rule 'ProdRule3' SET STGPPOOL DB-pool2 FOR FILESET (DB) where NAME LIKE '%logs%'
rule 'ProdRule4' SET STGPPOOL Data-pool1 FOR FILESET (Users)
rule 'ProdRule5' SET STGPPOOL Temp-pool where NAME LIKE '%temp%'
rule 'ProdRule6' SET STGPPOOL Temp-pool where NAME LIKE '%.log'

```

---

In our example, we have two File Placement policies defined, with only one, the `production_policy`, which is active. Each policy contains a list of file placement rules.

In this example, the first rule of the `production_policy` specifies that the directory “data” under fileset DB will be located into the storage pool “DB-pool1”. All files that don’t match any rule are created in the DEFAULT-POOL.

Figure 18-3 shows the SVC and SFS detailed configuration.

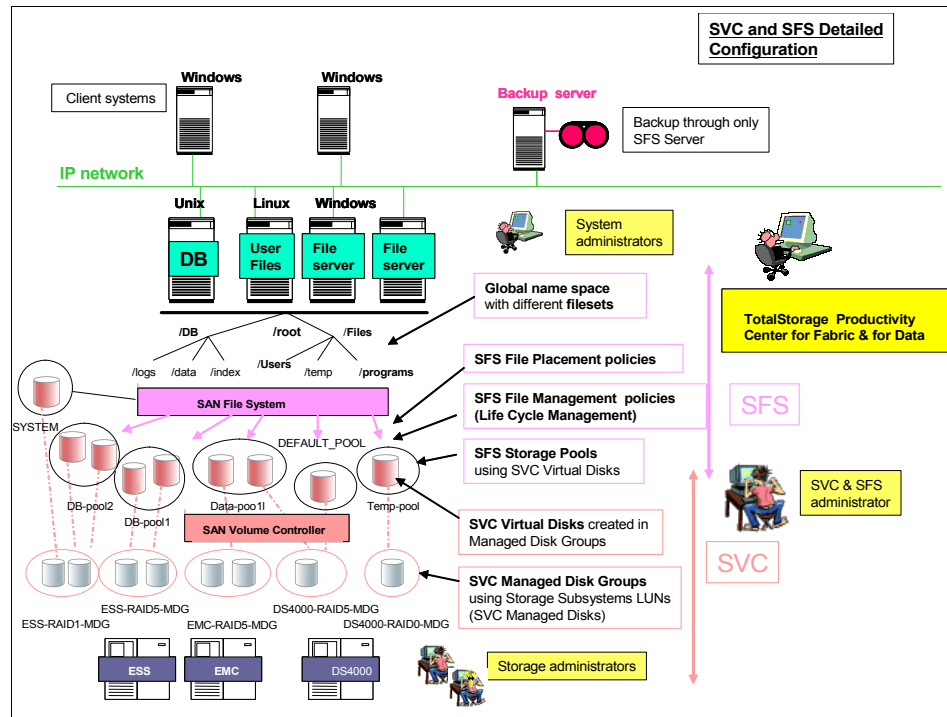


Figure 18-3 SVC and SFS detailed configuration

A Life Cycle Management (LCM) has been also implemented with the SFS policy based File Management feature: a File Management policy is also a set of rules, which specify conditions for either automatically move files from one storage pool to another, or for delete them from the storage pool.

Example 18-3 gives an example of a File Management policy LCM-Prodction.txt. In this example, two rules have been defined:

- ▶ The first rule will migrate files that are greater than 10 MB, from storage pool DEFAULT-POOL to storage pool Data-pool1.
- ▶ The second rule will delete all files that have not been accessed since the last 30 days from storage pool Temp-pool and from all filesets.

### Example 18-3 Life Cycle Management with File management policy

```
mds2:/home/LCM # cat LCM-Production.txt
RULE 'BigFiles' MIGRATE FROM POOL 'DEFAULT-POOL' TO POOL 'Data-pool1' FOR
FILESET ('Users') WHERE SIZE >= 10 MB
RULE 'OldTemps' DELETE FROM POOL 'Temp-pool1' WHERE AGE >= 30 days
```

With SFS, backup operations could be then simplified and performed now through a single agent, installed on any SFS client, because it has access to the storage space shared by all the servers. By reducing the number of backup agent licences, this configuration will directly reduce the storage infrastructure management cost.

### **IBM TotalStorage Productivity Center for Fabric and**

### **IBM TotalStorage Productivity Center for Data**

These two products have to be installed on the management workstations just like their corresponding agents, which may be installed on the servers connected to the SAN.

Each product is configured and customized to define the different reporting policies or to implement the automatization features such as the automatic file system extension. The following figures show some examples:

- Figure 18-4 shows the storage topology as deployed by IBM TotalStorage Productivity Center for Fabric.

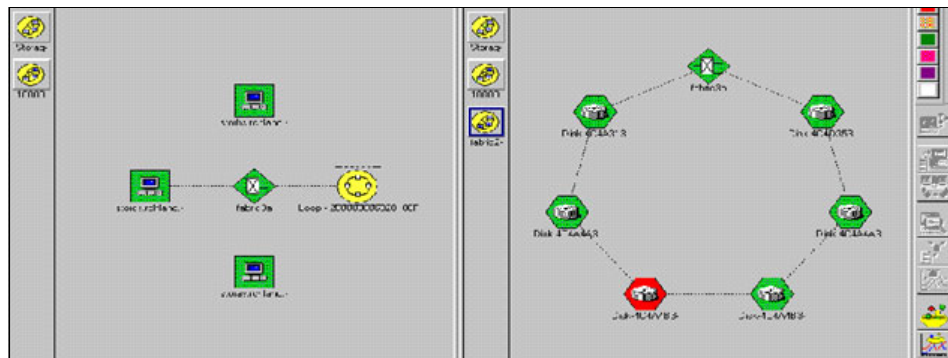


Figure 18-4 Visualization of the storage topology

- Figure 18-5 shows the storage identification as displayed by IBM TotalStorage Productivity Center for Data.

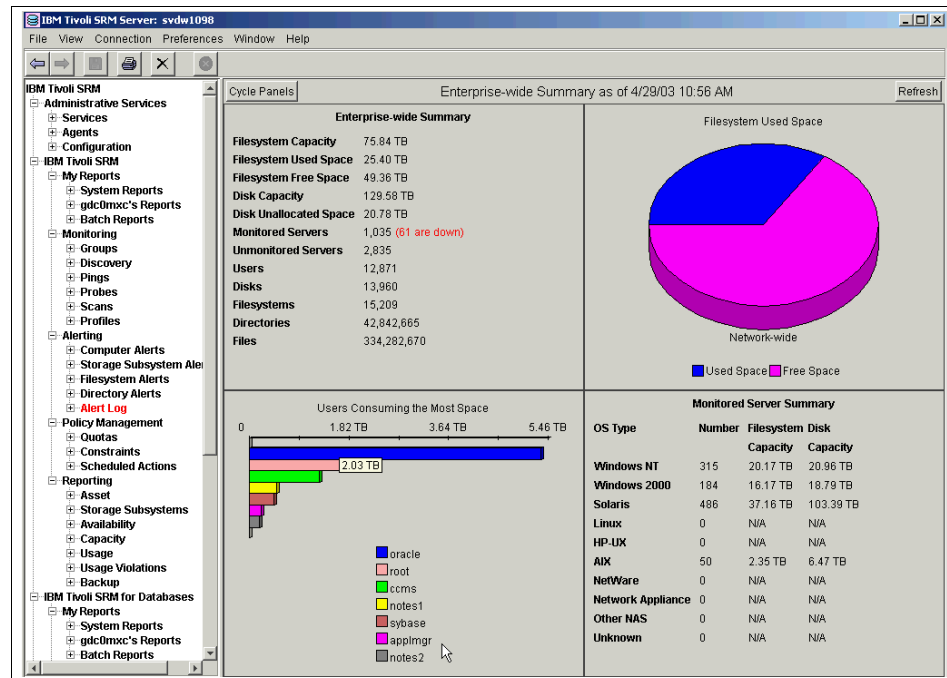


Figure 18-5 Storage identification

## Setting up the environment

Now, let's discuss the different tasks and processes necessary to integrate IBM Storage virtualization products and IBM storage management products into the existing infrastructure and implement our solution.

This new storage infrastructure implementation can be divided into six different main tasks which have to be performed in the given order. The first three tasks are related to the initial installation of the two storage virtualization products SVC and SFS. Then, the fourth and fifth tasks correspond to their management, and the last two tasks are for TotalStorage Productivity Center products installation and backup server configuration modification.



### **Task 1: SAN Volume Controller integration and configuration**

This task sets up the modifications required in the SAN and the storage subsystems:

- ▶ The SVC Cluster (SVC nodes and Master Console) has to be physically installed and connected to the existing storage infrastructure (FC fibers and Ethernet network).
- ▶ The fabric zoning configuration needs to be modified, according to the SVC requirements, so that the SVC nodes can get access to the storage subsystems that they will consolidate.
- ▶ The SVC basic configuration has to be done and a management console has to be configured to manage the SVC cluster, either by CLI or GUI.
- ▶ The free space on the different storage subsystems will be configured according to the target storage architecture; all the new created LUNs have to be mapped to the SVC nodes.

**Note:** To simplify the integration procedure, we suppose that there is enough free space available on the different storage subsystems to hold both the existing storage configurations and the new virtualized configuration we created and into which existing data will be migrated. When there is not enough free space available, it is necessary to back up the existing data to free some space into the storage subsystems. SVC itself provides an *Image Mode* migration feature that allows data LUNs to be converted directly in Virtual Disk; however, in our case, this interesting feature is not usable because Virtual Disks will not be used directly by application servers but rather through the SAN File System.

- ▶ Under the SVC, these LUNs, which are discovered as Managed Disks, have to be grouped into different Managed Disk Groups from which the Virtual Disks could be then created.

### **Task 2: SAN File System integration and configuration**

This task configures the SFS clients:

- ▶ The SFS Cluster (Meta Data servers) needs to be physically installed and connected to the existing storage infrastructure (through FC fibers or Ethernet network).
- ▶ The fabric zoning configuration needs to be modified, to reflect the SFS requirements, to offer an access to the SVC Cluster and to the different application servers.

- ▶ On the SVC, one or more virtual disks are mapped to the SFS Meta Data servers; they will be used in the SYSTEM storage pool to hold the SFS internal data base. The other virtual disks have to be mapped to the different application servers according to the SFS defined architecture.
- ▶ The SFS cluster needs to be created and its basic configuration has to be performed.
- ▶ The storage pools are then defined with the SVC virtual disks, which are seen as volumes under the SFS configuration.
- ▶ The Filesets are defined to reflect the Global Name space organization.
- ▶ As SFS clients are going to use SVC virtualized storage space, it is necessary to install the *Subsystem Device Driver* (SDD) software to deal correctly with the SAN multipathing and any failover situation that can occur (SVC nodes failover).
- ▶ The SAN FS Client software has to be installed on all the application servers that require an access to the SAN file system. When this installation is done, the SAN file system appears as a new file system on the UNIX or Linux clients or as a new drive letter on the Windows clients.
- ▶ The previously defined Fileset organization tree has to be visible but not yet accessible since file sharing has not been defined.
- ▶ This file sharing must now be configured carefully. There is the choice to implement either a “Basic Heterogeneous File Sharing” or an “Advanced Heterogeneous File Sharing”:
  - The Basic File Sharing is easy to implement, but has some limitations on sharing and security aspects.
  - The Advanced File Sharing improves the flexibility and the security of the sharing space but is more complex to implement. Moreover, it requires availability and configuration of additional products such as Active Directory, NIS, or LDAP.
- ▶ The File Placement policies have to be defined with their different placement rules, and one policy has to be activated before a client can start accessing and filling the SAN File System.

### ***Task 3: Data migration into the SAN file system***

This task integrates the existing data into the new virtualized storage architecture.

- ▶ From each application server accessing the SAN file system, migrate the data from the existing LUN based storage configuration into the new SAN file system tree. SFS provide a migration tool named *migratedata*, which can help with this process.

**Note:** As data is copied through the SAN file system tree, into the different storage pools, and by following the rules of the active File Placement policy, it is important to activate the correct policy before starting the migration.

- Once data has been fly-migrated into the new virtualized storage space, the storage subsystems LUNs, that were used by the old configuration, can be deleted; the corresponding space can now be integrated into the SVC global space as new Managed Disks.

#### **Task 4: Basic SVC and SFS Management**

This task describes the basic management tasks that have to be performed on the storage virtualization products, including usage of new features such as virtual disk migration tasks that are part of the day-to-day management:

- The first management task is to monitor how full the different storage spaces are becoming; the used space and free space need to be verified on the following components:
  - SVC Managed Disk Groups
  - SVC Managed Disks / SFS Volumes
  - SFS Storage Pool
  - SFS Filesets

**Note:** The SFS log files entries and different kinds of alert messages (trap SNMP, mails, etc.) also should be checked if you want to discover some possible errors or reached thresholds.

The following actions are now possible, at the SVC or the SFS level, to deal with problems discovered:

- If you need more space into one Managed Disk Group, you can increase its size by dynamically adding a new Managed Disk (LUN).
- If you want to redistribute some SFS Volumes (SVC Virtual Disks) between the storage pools, you can remove any volume from a storage pool, dynamically and without losing data; in fact, before removing the volume from the storage pool, SFS will automatically copy its contained files on the other volumes in the same storage pools by a process called *Volume Drain*. Once the volume is empty and removed from the storage pool, it is freed and ready to be used in any other storage pool.
- One of the major SVC features, called *Virtual Disks Migration*, gives you the possibility to move dynamically any SVC Virtual Disk between Managed Disk Groups, allowing a real transparent migration of data between different heterogeneous storage subsystems:

- As shown in Example 18-1 on page 281, you could migrate the 350 GB Virtual Disk “Default-VDisk0”, which is used as a volume in the SFS DEFAULT\_POOL, from “DS4000-RAID5-MDG” to “EMC-RAID5-MDG” Managed Disk Groups. This migration produces the SFS Default storage pool (with its files) to be dynamically migrated between the DS4000 subsystem to the EMC subsystem without any application down-time, which was not possible in the older non-virtualized storage architecture.

**Note:** You can only move virtual disks between Managed Disk groups built with the same *extent size*.

### **Task 5: Advanced SVC and SFS Management**

Other more advanced configuration or management tasks can be performed at the SVC or the SFS level, for example, file movements or copy services operations that might have to be performed:

- ▶ If necessary, files can be moved between SFS storage pools, adding some flexibility in storage space usage and allowing any file placement modification. Files can be moved manually, by an SFS **mvfile** command on individual or group of files.
- ▶ Files can also be moved automatically by the new SFS Life Cycle Management (LCM) feature. File Movement Policies have to be defined as in Example 18-3 and scheduled to run automatically at fixed intervals.
- ▶ The SFS administrator can create FlashCopy images of any SFS fileset. These images correspond to a fixed “picture” of the fileset and all its files; they are in the state they had at the time the FlashCopy was performed; they are saved into a special directory named `flashcopy` which is located at the root of each fileset; they are directly accessible by all the SFS clients of the SFS file system involved in the process.
- ▶ The SVC provides also a FlashCopy feature which is working at the Virtual Disks level; in this case, we can create FlashCopy components between a source virtual disks and a target virtual disk (with the same size); the target virtual disk, which is an image of the source virtual disk, could be used to improve backup operations.
- ▶ The SVC provides also a Peer-to-Peer Remote Copy (PPRC) feature, which can be established between virtual disks of the same size. PPRC linkages can be established inside the same cluster (intra-cluster relations) or between two different SVC clusters (inter-cluster relations).
  - Inter-cluster links are used to build disaster tolerant infrastructures between two physical sites, each one having an SVC Cluster.

- Intra-cluster links can be used, to create clones of virtual disks that can be used for specific operations; for example to quickly build a test environment using a copy of the production data.

### ***Task 6: IBM TotalStorage Productivity Center installation***

Once the SVC and SFS clusters are installed and the data migrated to their virtualized storage space, the TotalStorage Productivity Center products can be installed and configured:

- ▶ For each product, the manager part of the software has to be installed on a dedicated management workstation.
- ▶ The agent software has to be installed on the different application servers connected to the storage space:
  - The TotalStorage Productivity Center for Fabric agents act as Inband agents (through the SAN) in the Fabric topology discovery and reporting mechanisms (by opposition to the outband agents that are contacted through the network).
  - The TotalStorage Productivity Center for Data agents collect data information from the managed system on which it is installed. They then send this collected data to the Manager; consolidated reports on storage usage can then be generated.
- ▶ Each product can be configured and customized to define in detail the reports or the graphs we want to generate.
- ▶ The Storage Management policies can also be defined to specify different kinds of alerts, quotas, and constraints that we want to monitor with the resulting notification actions.
- ▶ More advanced feature such as *Automatic File System Extension* can also be configured; it helps improve the automatic storage management level.

### ***Task 7: Backup server configuration***

Modifications may be needed on the backup server to reflect the new storage architecture. Changes will depend on how the backup is implemented: this optional task is not detailed here.

## **18.4.5 Benefits and summary**

Once fully deployed, the IBM's Storage Virtualization and Infrastructure Management solution builds a consolidated storage solution, with flexibility and scalability; this solution makes easier storage management and reorganization procedures to meet business changes.

The final users, such as the storage administrators, now can enjoy the following benefits with this new storage infrastructure:

- ▶ The storage capacity can be always maintained at an optimum level to avoid any lack of storage resource, by using the dynamic storage reorganization capability.
- ▶ The utilization of under-utilized storage resources is improved without requiring an investment in additional capacity, by dynamically allocating capacity to applications that require it.
- ▶ New applications can be more quickly deployed.
- ▶ IT storage infrastructure can be manipulated in real time by using defined business policies.
- ▶ Data sharing between different operating systems is real, easier, and faster.
- ▶ Data migration between all the different available storage subsystems is possible with no application disruption.
- ▶ Storage management is simplified and less expensive due to the centralization management process.

## **18.5 Product positioning: Conclusion**

In the following sections we describe how the related storage products and features bring a solution to the global On Demand Operating Environment.

### **18.5.1 IBM TotalStorage SAN Volume Controller**

SVC allows an aggregation of all the SAN storage space and the control from a central management point.

This provides a more flexible utilization of the different existing storage subsystems. It is even possible, with SVC, to migrate dynamically and without any impact on the application traffic, virtual disks between heterogeneous storage subsystems. That allows you to satisfy more quickly, any infrastructure change requests or maintenance operations on a storage subsystem.

SVC provides also advanced copy features, such as FlashCopy and PPRC, which allow the creation of images, clones, or replicas of managed disks which all can help to reduce the response time in storage management operation.

## 18.5.2 IBM TotalStorage SAN File System

This product, with its policy-based file-placement feature, also allows a better storage space utilization through many customized rules.

File location can be also controlled and modified, either manually or through a policy-based file-management feature which acts as a Life Cycle Management feature when it is automatically scheduled.

SFS allows also the storage pools to be increased or decreased by adding or removing volumes dynamically. Volumes can be removed from a storage pool even if they hold some files. The SFS Volume drain feature can automatically copy these files on the other volume before removing the volume from the storage pool and allowing its immediate usage for another storage pool.

SFS allows also the creation of FlashCopy images at a fileset level, which is directly accessible by SFS users to retrieve for themselves, old versions of their files.

All of these features combine to improve the capacity of the infrastructure to become more of an On Demand Operating Environment.

## 18.5.3 IBM TotalStorage Productivity Center

The IBM TotalStorage Productivity Center product, in addition to its centralized management capability, also provides some policy-based advanced features such as the automatic file system extension, which contribute to the overall on demand capability of the solution.

## 18.6 Summary

In this chapter we have described how a business requirement for a storage infrastructure consolidation and optimization with management simplification has been addressed by the products available today. A scenario describing how these products have been used together to provide a solution has been detailed, and a step-by-step implementation procedure has been documented.

In this chapter, we listed how the products and the global solution fit into the On Demand Operating Environment framework that allows customers to evolve their environment in a stepwise fashion.







## How to monitor using EWLM

This chapter focuses on the need for efficient monitoring in today enterprise to ensure that work requests are performing as expected in the more and more common multi-tiered environment. One of the most important elements of successful performance management is understanding transaction response times and topologies. So, no matter how complex your business environment, you are likely interested in answering the following questions:

- ▶ Are work requests completing successfully? If not, where are they failing?
- ▶ Are successful work requests completing within the expected response time? If not, where are the bottlenecks?
- ▶ How many work requests were completed over some period of time compared to prior periods? Is the workload growing?
- ▶ Are system-level resources being used for optimal performance? If not, can they be dynamically redirected to alleviate bottlenecks?

This chapter provides you accurate answers to these questions by giving you the ability to do the following:

- ▶ Identify work requests based on business priority.
- ▶ Track the performance of work requests across server and subsystem boundaries.

## 19.1 Introduction

Monitoring in an On Demand Operating Environment is the capability to understand the topology and control the performance of the end-to-end business applications.

EWLM allows you to define business-oriented performance goals for an entire domain of servers, and then provides an end-to-end view of actual performance of the applications relative to those goals. EWLM is a key component of the IBM Virtualization Engine. The Virtualization Engine solution simplifies management and utilization of heterogeneous IT resources. As a part of this solution, EWLM offers the following benefits:

- ▶ Provides an understanding of the actual flow of work, based on dynamic discovery.
- ▶ Automatically detects server and application topologies so that it can rapidly and efficiently identify the likely origin of performance problems.
- ▶ Allows users to define performance goals in policies that are similar to service level agreements.
- ▶ Provides load-balancing capabilities to help ensure that performance goals are met.

## 19.2 General strategy

The components needed to implement this approach have to respect the general attributes of an On Demand Operating Environment. Specifically, they must be:

- ▶ **Self-managing:** One element of the EWLM solution is the Control Center, where EWLM has the capability to keep track of the different servers joining and leaving the management server. Also, there is no need to define where and which kind of applications are running on the different servers belonging to the management domain: EWLM will be able to dynamically discover the application multi-tier topology eliminating any administrative overhead.
- ▶ **Scalable:** The solution is capable to adjust to a growth of the enterprise, in terms of servers or applications.
- ▶ **Resilient:** The security components are integrated parts in the monitoring solutions. EWLM supports stateful inspections, firewalls, HTTP Proxy, and SOCKS.
- ▶ **Economical:** Any solution today needs to be economical, and a monitoring solution is no different. A centralized solution that addresses multiple platforms and takes advantage of the unique facilities provided by some of those platforms takes advantage of economies of scale. In addition, it reduces the administrative overhead.

- **Open standards-based:** The only way that a centralized solution can provide consistent monitoring to a large number of platforms, operating systems, applications, and data stores is by building on and using open standards, such as ARM V4.

## 19.3 Solution components

EWLM is a product in the IBM Virtualization Engine Suite for Servers that can dynamically monitor and manage distributed heterogeneous workloads to achieve user-defined business goals.

EWLM *collects and aggregates* performance statistics so that you can compare a high-level goal to the actual performance at all times. The EWLM Control Center provides a user-friendly way of viewing this information. EWLM captures relevant performance statistics on each operating system instance within a management domain. It then reports an aggregated view of the real-time state of the domain to the EWLM domain manager. In the EWLM Control Center you can view detailed statistics that explain why the actual results were achieved, as opposed to the desired results. You can capture performance statistics in a tabular form for offline analysis or for integration into other performance reporting tools

In order for EWLM to report on these transactions and potentially manage the domain based on these statistics, *Application Response Measurement 4.0 (ARM)* has been implemented for application instrumentation. ARM is an Open Group standard composed of a set of APIs that can be used to collect response time information for enterprise applications.

Instrumentation provides the data to create the *topology of a transaction*, that is, what server instances or application instances a transaction flows through. In a typical Web transaction, for example “Buy Books,” the transaction starts at a Web server (edge server), flows to an application server, then to a database server. Without instrumentation, the transaction would look like three separate processes with three separate response times rather than one transaction. This is key to EWLM, providing end-to-end topology views and statistics for business transactions.

So, in addition to the EWLM product, it is also necessary that the middleware where the business applications are running is ARM instrumented. Today, ARM 4.0 supported middleware includes the following:

- DB2 UDB V8.2
- WebSphere V5.1.1
- WebSphere Plug-ins delivered with 5.1.1
- Independent IHS/Apache Web server ARM plug-ins
- Independent IIS Web server ARM plug-in

## 19.4 Scenario

So far we have discussed the EWLM strategy and different components that are part of the monitoring solution. Now, we will apply this environment to a simple On Demand Operating Environment scenarios for an hypothetical organization with a typical set of requirements.

### 19.4.1 Current environment

Let's take as a sample a fictitious environment, MyCo, Inc. for example, for a stock brokerage and banking Web site where customers can transfer funds, perform account queries, buy and sell stock, and handle mortgage loans. In this environment, Service Level Agreement (SLA) are in place to rule these business applications. This scenario describes how a domain policy was created for the banking applications.

#### Data centers

Figure 19-1 shows a simplified configuration layout for the Web applications:

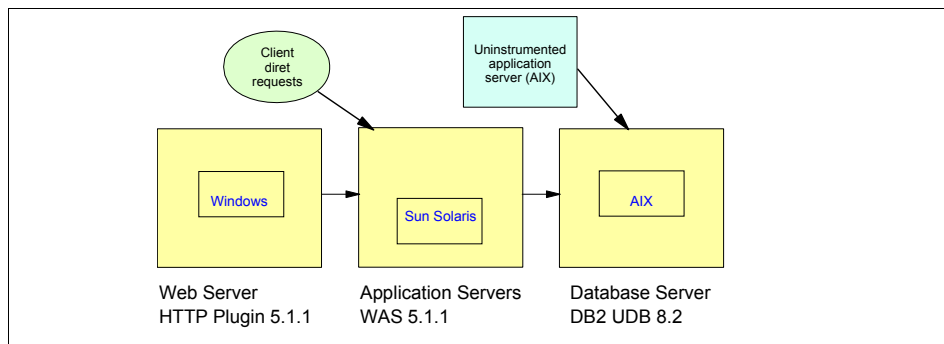


Figure 19-1 Web applications layout

- ▶ The EWLM managed servers are:
  - AIX 5L
  - Microsoft Windows 2003 Server
  - Sun Solaris 8
- ▶ The applications are running in the following environment:
  - Microsoft Internet Information Services (IIS) plug-in to WebSphere, with ARM-instrumentation enabled.
  - WebSphere Application Server 5.1.1, with ARM-instrumentation enabled.
  - IBM DB2 Universal Database Version 8.2, with ARM-instrumentation enabled.
  - Other Web application server, not ARM-instrumented.

## 19.4.2 Business objective

In this scenario, MyCo, Inc. wants to be able to assign a business goal to their end-to-end Web Banking applications transactions, and then been able to monitor them and report them according to the established SLAs.

MyCo's Web banking applications include a funds transfer, an account query, a mortgage rate advisor, stock buying for brokers, and a bank statement processing application.

The Vice President of finance requires quarterly reports on the performance of each of the bank region's brokerage buyers. They serve the northeast and northwest regions of the United States.

## 19.4.3 Technical objectives

MyCo, Inc. wants to get more from their existing IT infrastructures. They need to optimize manually-intensive performance analysis, specially during failure time to avoid outages in their service time. The following summarizes the difficulties and concerns that the company is facing today:

- ▶ Difficulties in monitoring the end-to end applications service levels:
  - Without effective tools to manage the complex heterogeneous multi-tiered environment, more personnel are required to monitor the infrastructure to ensure service level agreements (SLAs) are met.
  - Reporting on SLA attainment is a time-consuming and labor-intensive process. It requires manual comparison between data from multiple IT infrastructure components, and the customer-agreed SLA.
  - Traditional system management tools only focus on servers, routers, applications, disk space, and so on. These types of reports do not show which business aspects are affected if any of the components is down or unavailable.
- ▶ Elongated outage time because of the difficulties during the diagnosis of a performance problem in a such complex environment.
- ▶ Increasing in IT infrastructures to compensate the difficulties in balancing the workload and understanding the correct resources consumption

## 19.5 Scenario implementation

In the following sections, we introduce some simple concepts that build the monitoring solution provided by EWLM. Note that EWLM is an implementation of *policy-based performance management*.

## 19.5.1 Architecture components

Figure 19-2 describes the different components in an EWLM solution.




Figure 19-2 EWLM architecture

- The management domain:

The scope of management is a set of servers that you logically group into what is called an EWLM *management domain*. The set of servers included in the management domain have some type of relationship; for example, the set of servers supporting a particular line of business. The line of business may consist of multiple business processes, spread across a few servers or a thousand servers.

- The domain manager:

There is a management focal point for each EWLM management domain, called the EWLM *domain manager*. As shown in Figure 19-2, there is a one-to-one relationship: one domain manager instance per management domain. The domain manager coordinates policy actions across the servers, tracks the state of those servers, and accumulates performance statistics on behalf of the domain.

► The managed server:

On each server (operating system) instance in the management domain there is a thin layer of EWLM logic installed called the EWLM *managed server*. From one perspective the managed server layer is positioned between applications and the operating system. The managed server layer understands each of the supported operating systems, gathering resource usage and delay statistics known to the operating system.

A second role of the managed server layer is to gather relevant transaction-related statistics from middleware applications. The application middleware implementations, such as WebSphere Application Server, understand when a piece of work starts and stops, and the middleware understands when a piece of work has been routed to another server for processing, for example, when a Web server routes a servlet request to a WebSphere Application Server.

The managed server layer dynamically constructs a server-level view describing relationships between transaction segments known by the applications with resource consumption data known by the operating system. A summary of this information is periodically sent to the domain manager, where the information is gathered together from all the servers in the management domain to form a global view.

An important aspect of the EWLM approach is that all data collection and aggregation activities are driven by a common service level policy, called the EWLM *Domain Policy*. This policy is built by an administrator to describe the various business processes that the domain supports and the performance objectives for each process. For example, a book store would probably have a business transaction called “Add to shopping cart” that would have a performance goal of perhaps a 1.5 second response time.

To permit the EWLM domain manager to construct an end-to-end view of each type of business transaction, the processing segments performed by each participating middleware instance must be correlated, meaning pieced together. For example, the Add to shopping cart transaction would be received by a Web server, it could flow to a WebSphere Application Server instance which might update a database. The transaction might even check your credit limit by routing a sub-transaction to another server. The tricky part of this correlation is to get all of the applications in the path of each transaction to cooperate, without knowing that they are cooperating.

► The control center:

The final aspect of EWLM to introduce is the EWLM *Control Center*, a browser-based application server tailored to the needs of an EWLM administrator, analyst, and operator. This is where all the EWLM concepts come together; where you can create a service level domain policy and activate that policy on hundreds of servers with a single click. Reporting data

is then available to let you view performance from a business perspective. And if you need the details, you can see the topology of servers and applications supporting each transaction type and understand where the time was spent.

The information is organized in such a way that an administrator or analyst can easily drill-down to what is relevant.

## 19.5.2 Implementation of IBM EWLM

In order to achieve the business objective of the end-to-end business transactions reporting, the following activity tasks list has been built:

- ▶ Identify and install the EWLM component and define the Control Center element
- ▶ Define service classes for the different applications offered on their Web site.
- ▶ Specify which work should be assigned to the service classes.
- ▶ Report the performance of brokerage buyers for the northeast and northwest regions.
- ▶ Include the application running on a competitive Web application server.
- ▶ Use the above to define a domain policy representing the Web banking environment.

Let us consider in detail what needs to be done to accomplish these activities. The following paragraphs guide you step by step through a scenario where you will install the EWLM component first and then define an EWLM domain policy for MyCo's bank Web site.

### Step 1: Setting up the environment

EWLM is an optionally installable system service of the IBM Virtualization Engine. You can install EWLM Domain Manager, EWLM Managed Servers for IBM and non-IBM platforms, or some combination of these. The EWLM Control Center is installed as part of the Domain Manager installation.

Once you complete the EWLM installation, you need to run a set of configuration scripts to get the Domain Manager, EWLM Control Center, and Managed Servers up and running.

#### ***Installation of EWLM code***

The installation process for EWLM is part of the Virtualization Engine installation wizard. There are two phases in this installation process:

- ▶ The first phase involves transferring the media, during which the installation wizard copies images of the CDs to whichever destination you select.



- ▶ The second phase is the installation of the Virtualization Engine (VE) code and the VE services.

The installation wizard allows you to select which VE system services to install and performs the actual installation. In the case of the EWLM installation, the selection of the EWLM system services will install the Domain Manager and the EWLM Control Center as well as the required runtime environment. In addition, the Managed Server installation images are copied to the main server. These should be later distributed and installed to the appropriate platforms as a separate step.

The EWLM installation requires you to go through the following steps:

- ▶ Ensure that your system meets the minimum hardware and software requirements.
- ▶ Complete a planning checklist for your platform. The checklist is included in the installation process as part of the installation assistant.

When the installation wizard completes, you have all the EWLM code install on the server. Now you need to complete the installation with the following post-installation tasks: configure the domain manager, start the WebSphere Application Server instance on the domain manager, configure the managed server, and enable ARM services on the operating system.

### ***Configuring the domain manager***

The domain manager is the central point in the management domain where all the performance information is collected and aggregated.

The configuration of the domain manager requires you to run a script or an installation wizard to define which is the IP address or hostname and the port that the domain manager will use to communicate with the managed servers that will be part of this management domain.

### ***Configuring the managed servers***

When you install the EWLM code on the domain manager, you are also given the option to install the images for the managed servers. This means that the code for the managed server is copied to the domain manager but it needs to be pushed out to the managed servers. You need to manually transfer the EWLM managed server code from the domain manager to each managed server in the management domain. This distribution can be done either by file transfer protocol (FTP) using binary mode on AIX and Solaris, or by mapping a drive and copying the files to the file system, or by any other distribution mechanism that works in your environment.

Once you have deployed the code on the managed servers, you need to run a script or an installation wizard to configure the managed server domain manager: this will define which is the IP address or hostname and the port that the domain manager will use to communicate with the managed servers that will be part of this management domain.

## **Step 2: Define performance objectives in service classes**

After having installed and having configured the EWLM domain manager and the EWLM Control Center is started, the following scenario is a sample of the tasks that are required to define your applications to EWLM through the definition of the domain policy; EWLM will then be able to start monitoring the environment.

For example, we can use the domain policy name Sample Banking Domain Policy and go through the steps required to complete the definition and activation of the policy in the management domain:

### ***Define the business applications for the Web banking environment***

We first need to identify which transaction application components we want to monitor and understand their environment:

- ▶ Whether the work is transaction-based or process based
- ▶ On which middleware the business applications run
- ▶ Which of the middleware applications are ARM 4.0 instrumented

In this example, we assume that the following functions need to be included:

- ▶ Funds transfer
- ▶ Account query
- ▶ Mortgage rate advisor
- ▶ Brokerage stock buying
- ▶ Static Web page content that includes graphics and text
- ▶ WebSphere Application Server EJB application to process bank statements

All are transactions based work and all run on the ARM 4.0 instrumented middleware applications except for the mortgage rate advisor application, which runs on the Web application server from the bank merger (herein known as the non-instrumented Web application server).

- ▶ The non-instrumented Web application server does, however, pass requests to DB2 UDB, which is ARM 4.0 instrumented. This means that the non-instrumented Web application server can be included in the end-to-end domain policy only as process-oriented work because transaction-level performance data is not available for a non-instrumented application. Because the non-instrumented Web application server passes data to DB2 UDB, which is instrumented, the UDB portion of transactions can be monitored and managed by EWLM and therefore, can be included in the domain policy.

### ***Define the performance objectives for the work***

Service classes based on the *performance objectives* for the Web banking work need to be defined and a *business importance* needs to be assigned to the work.

Importance is a value that defines how important it is to the business that the performance objective be met. The performance objectives depend on whether the work is considered transaction-based or process-based. Response time objectives usually apply to transaction-based work. For process-based work, EWLM provides a goal type called velocity. Velocity goals are appropriate for long-running work such as system daemons, batch jobs, and server processes. Velocity defines how fast work should run when ready, without being delayed for processor, storage, and I/O.

So, the next step is to assign *service classes* to the different components, with response times defined in their service level agreement as illustrated in Table 19-1. Additionally, other service classes can be defined to be able to monitor additional workloads. Once we have the list of the applications we want to monitor with the business goal, we are ready to map this into an EWLM domain policy.

*Table 19-1 Service classes*

<b>Service class</b>	<b>Importance</b>	<b>Performance goal</b>
WebSphere EJB Service Class	Highest	85% complete in 4 seconds
Account Queries	High	80% complete in 3 seconds
Brokerage Buyers	Highest	80% complete in 1 seconds
Stock Quotes	Medium	80% complete in 10 seconds
Static Web Serving	Low	90% complete in 20 seconds
Mortgage Rates	Medium	1 minute average response time
Non-Instrumented Web Applications	High	Fast Velocity
Web Browsers	Medium	Moderate Velocity

For the mortgage rate advisor application, the end-to-end response time and the time the transactions normally spend in the database query should be considered. Because the Web application server is not instrumented, the performance objectives defined in the domain policy represent only the portion of time the transactions spend in DB2 UDB.

The application is not quite real-time in the sense that the Web site provides a wizard to query mortgage rates with the rates sent as an e-mail, so the turnaround need not be instantaneous. The service level agreement states that an e-mail will be sent within the next hour after a query is made; the portion of the time spent in DB2 for the query is estimated to be about 20 seconds, so the response time objective is set for 1 minute.

For the mortgage rate advisor application, EWLM can manage and monitor the non-instrumented Web application server regions. To do this, a service class for the non-instrumented Web application server region with a velocity goal needs to be defined. A general name for this service class should be used so if another bank merger occurs later, and they would like to integrate a non-instrumented application, it can be included in this service class. Velocity is specified as a category: Fastest, Fast, Moderate, Slow, or Slowest. The processes running on Web Browsers in the banks can also be managed. The browsers, Internet Explorer and Netscape Navigator are processes, and are also assigned a velocity goal.

For work not otherwise categorized, EWLM provides a service class in every domain policy called *EWLM Service Class*. It is assigned a discretionary goal. Discretionary goals specifies that work is to be completed when resources are available. A discretionary goal type can be assigned to a service class for low priority work that does not require a particular performance goal. EWLM processes discretionary work using resources not required to meet the goals of other service classes.

### ***Define the workloads***

Based on the service classes, the following workloads can be defined for reporting purposes:

- ▶ WebSphere Applications, for the WebSphere EJB application that processes the bank statements.
- ▶ DB2 Applications, for the mortgage advisor transactions originating from the non-instrumented application running in DB2 UDB.
- ▶ Non-Instrumented Applications, for the Web application server process and the Web browser processes.
- ▶ Web Site Applications, for the funds transfer, account query, mortgage rate advisor, and brokerage buyer transactions, and the static Web serving work.
- ▶ EWLM provides a workload, EWLM Workload, which is assigned to the EWLM service class.

### Step 3: Define work to service classes

After the Web banking service classes are defined, the next step is to define which work should be assigned to each service class by using transaction classes and process classes. The transaction classes define the transactions that should be assigned to a service class. The process classes define the processes that should be assigned to a service class.

The transaction class or process class definitions contain the rules to identify the work. The attributes of incoming work are compared to these rules and, if there is a match, the rule is used to assign a transaction or process class to the work. For example, a transaction class may contain a rule that specifies that all transactions from server name `www.mycobank.com` should be assigned to the Static Web Serving service class. The attributes of the work are called filters, and include URI, system name, hostname, user name, and so on. EWLM provides some general filters, and the applications or platforms provide some filters specific to their environment, such as EJB name for WebSphere or Executable Path for AIX.

EWLM monitors and manages resources to meet the goals defined for service classes, and monitors based on transaction classes and process classes. You can view performance data on transaction classes and process classes in addition to performance data on service classes. You can use transaction classes and process classes to not only define which work should be assigned to a service class, but to define which work within a service class you would like to specifically monitor.

With the service classes defined, the next step is to define transaction classes and process classes.

#### ***Define the entry applications for the service classes***

A service class definition is required in order to be able to define a transaction classes. The first ARM-instrumented application that processes the transactions representing that service class in the Web Banking environment needs to be determined. The first ARM-instrumented application that processes the transactions specifies the start of the end-to-end performance for the service class. The first application is known as the *entry application*. The entry application or edge is where EWLM looks up the classification rules defined in the transaction classes.

For each of the service classes representing transaction-based work, the entry application needs to be determined as described in Table 19-2.

Table 19-2 Edge applications

Service class	Entry application
Static Web Serving	IIS plug-in to WebSphere
Account Queries	IIS plug-in to WebSphere
Brokerage Buyers	IIS plug-in to WebSphere
Stock Quotes	IIS plug-in to WebSphere
WebSphere EJB Service Class	WebSphere Application Server
Mortgage Rates	IBM DB2 Universal Database

- ▶ The entry application for the Static Web Serving, Account Queries, Brokerage Buyers, and Stock Quotes service classes is IIS, which, in EWLM, is included in the general IBM Web serving plug-ins.
- ▶ The entry application for the WebSphere EJB service class that processes the bank statements is WebSphere Application Server, since the EJB is from a client direct request.
- ▶ The entry application for Mortgage Rates service class is IBM DB2 Universal Database. While the transaction originates in the Web application server from the bank merger, the Web application server is not ARM-instrumented, so it cannot be the entry application.

Then, though the EWLM Control Center, the entry applications are added to the domain policy.

### ***Define the transaction classes***

MyCo Inc. uses a very strict naming convention for its systems and transactions, which is in fact very useful for defining the classification rules. The following tables identifies the applications that EWLM will monitor. In addition, it specifies the details of the domain policy for each application.

The IBM Web serving plug-in will be used as the entry application for most of MyCo's 3-tiered banking applications. Table 19-3 shows the list of the applications with the correspondent rules able to classify this incoming workload to the specified service class.

Table 19-3 Transaction classes and correspondent classification rules

Transaction Class	Description	Service Class	Rules
Static Web Serving	Pages and graphics for the bank Web site	Static Web Serving	HostInfo Equal www.mycobank.com
Stock Query	Stock ID passed as DB2 query	Stock Quotes	QueryString Equal WBstock="stockID"
NorthWest	NW region brokerage buyers	Stock Quotes	EWLM:System Name Equal myco.northwest.com
NorthEast	NE region brokerage buyers	Stock Quotes	EWLM:System Name Equal myco.northeast.com
Account Query	Account query transactions to UDB	Account Queries	QueryString Equal WBname="username% " &WBaccount="accountID%"
Default - IBM Web serving plug-in	Default transaction class for IIS	EWLM Service Class	(\*) Equal (\*)
Default - WebSphere	EJBs supporting the banking applications	WebSphere EJB Service Class	(\*) stringMatch (\*)
Default - IBM DB2 Universal Database	Default transaction class for DB2	Mortgage Rates	(\*) stringMatch (\*)

Because the Vice President of finance wants to have reporting on the performance of the northeast and northwest brokerage regions, the IIS server names include the region, so this workload can be classified and monitored.

### ***Define the platforms for the process***

In addition to the transactions, also process works we want to monitor need to be defined to the domain policy. Process-based works almost always run on a single server or a set of servers on a single platform. To define the work in process-based service classes, the first step is to define the platform on which the processes run, as shown in Table 19-4.

Table 19-4 Platforms

Service class	Platform
Non-Instrumented Web Applications	AIX
Web Browsers	Windows

This is because the non-instrumented Web application server runs on AIX and the Web Browsers run on a Windows platform, in this example.

### ***Define the process classes for the process based service classes***

Now that the platforms have been defined, we need to classify the work and associate it with a process class. Table 19-5 shows a summary for our sample scenario.

Table 19-5 Process classes

Platform	Process Class	Description	Service Class	Rules
AIX Platform	Web Application Server Processes	Example process class for uninstrumented applications	Non-Instrumented Web Applications	ExecutablePath Equal webapserver/webpages
Windows	Netscape Web Browser	Netscape Navigator	Web Browsers	ExecutablePath Equal C:\ProgramFiles\Netscape\Communicator\Program\netscape.exe
Windows	IE Web Browser	Internet Explorer	Web Browsers	ExecutablePath Equal C:\Program Files\Internet Explorer\IEXPLORE.EXE

## **Step 4: Deploy the domain policy**

Now that all the Web Banking environment has been classified and has performance goal assigned in the domain policy, we can deploy it to make it available to the domain and then start monitoring.

### ***Deploy the domain policy and activate the service policy***

The deployment of the domain policy and the activation of the Banking 2004 Service Policy is an easy task.

From the EWLM Control Center, it is possible in fact to select and deploy the domain policy across all the servers participating in the management domain.



During the deployment, it is also possible to activate immediately the policy. This will enable all the performance definitions in all the servers in the management domain. It is now possible to begin to use the Monitor section of the EWLM Control Center to view the performance of the work requests in your EWLM domain.

## **Step 5: Monitoring performance**

Now that the domain policy has been deployed and the service policy is activated, you can start looking at the performance of your applications through the Control Center. The Monitor section of the EWLM Control Center provides many views to obtain performance data.

The most common monitoring functions are as follows:

- ▶ Service classes that meet or do not meet their performance goals.
- ▶ Topology data that shows the specific applications and servers that are processing a particular service class, transaction class, or process class.
- ▶ Transaction statistics for a specific time interval.
- ▶ Performance index for a specific time interval.
- ▶ Applications that are classifying work requests to service classes. These applications are referred to as entry applications.
- ▶ Managed servers that are processing work requests for a particular service class or, more specifically, a process class. In addition, you can view the percent of the managed server's processors that are processing work requests that EWLM is monitoring.
- ▶ Transaction statistics for a service class, transaction class, or process class that include the number of transactions that have completed successfully or unsuccessfully.
- ▶ Server and application contributions to the end-to-end performance of the class (service, transaction, or process) such as response time, transaction counts, and delay percentages.

On the EWLM Control Center, there are six high level monitor views which can be used to start analyzing the performance of the workload.

- ▶ High level views:
  - Exceptions
  - Workload
  - Service classes
  - Transaction classes
  - Process classes
  - Managed servers

From each of the high level views listed above, selectable reports can be obtained by choosing a pull-down option. The more detailed information — or low level views — includes:

- ▶ Detailed monitor reports on the classes and managed servers
- ▶ Topology reports for the application and the servers with further detailed statistics
- ▶ Real time performance monitors:
  - Goal achievement
  - Processor utilization
  - Transaction count and rate

The first step in determining how well the performance is in your EWLM domain is to read and interpret the data from the high level monitor views. These views provide you with high-level performance data and indicate what kind of performance you are achieving in your environment, as described in Table 19-6.

*Table 19-6 High level monitor reports*

Exception reports	Specifies the service classes that are not meeting their performance objectives. Use this report to view performance statistics that are specific to service classes that are not meeting their goals. A service class is listed in the exception report if its performance index (PI) is greater than 1. The higher the PI, the greater variance of the service class achieving its goal. The PI should be 1.0 or less. If no services classes are listed in this report, all of the service classes are meeting or exceeding their goals.
Workloads	Identifies the workloads that are defined in the deployed domain policy. Use this report to view whether a group of service classes (a workload) are meeting their goals or not. In addition, if a service class appears in the Exceptions report, you can use the Workloads view to determine which workload is affected the service class that is not meeting its goal.
Service classes	Specifies all of the service classes that are defined in the deployed domain policy. You can view performance statistics for each service class.

Transaction classes	Specifies all of the transaction classes for each application defined in the deployed domain policy. You can view performance statistics for each transaction class. This view is particularly useful if a service class is not meeting its goal. You can use this view to pinpoint the transaction class within the service class that can cause the service class to miss its goal. You can view the average response times for each transaction class in a particular service class to determine which are missing their goals.
Process classes	Specifies all of the process classes for each platform defined in the deployed domain policy. You can view performance statistics for each process class. This view is similar to the Transaction classes view but the work requests included in a process class are initiated by a server process rather than an application.
Managed servers	Identifies all of the managed servers in the EWLM domain. You can view the status of each server and view data specific to each server. The managed server details include the services classes that are processed by the managed server. This is particularly useful if a service class is missing its goal due to the managed server not having enough resources. You can view the managed server's details to determine if the processor usage is acceptable for the number of transactions that it is processing.

For each of the six high level monitor views, there is more detailed information that you can obtain. This lower level performance data is critical in diagnosing workload problems running in a heterogeneous environment. From the EWLM Control Center, we can obtain performance data across multiple platforms. We can obtain information such as CPU, response times, processor delay and queue times. We can also get a visual of the middleware and servers that a transaction traverses from execution start to finish. We can narrow the scope to where problems may be occurring, such as communication problems or excessive wait times. Table 19-7 describes what you can provide from the low level monitor reports.

Table 19-7 Low level monitor reports

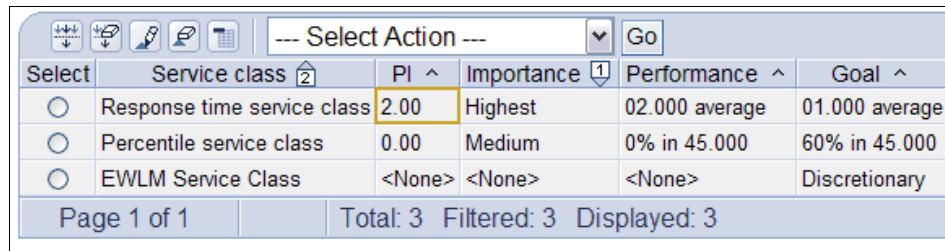
Details reports	Specifies transaction and performance details for a specific report. The specific data contained in each details report depends on what type of report you are viewing. For example, you can view a Service class details report, Process class details report, or Managed server details report. Use these details to help you determine whether a problem exists with a specific managed server or a problem exists with work associated with a specific process or service class.
Topologies	Provides a graphical view of the managed server or applications in your EWLM domain. You can view an application topology or managed server topology. From the topology view, you can determine which servers are processing work for each application in the EWLM domain. In addition, you can view the amount of resources that each application or server is absorbing to process work requests that are processed in the EWLM domain. For each topology, you can also view more granular data such as the average response times, number of delays, number transactions that completed successfully or unsuccessfully, and many other granular statistics. You can use all of this data to determine if there is a performance problem in your EWLM domain.
Real-time performance monitors	Provides graphical views that EWLM updates automatically. Use these graphs to determine how work is processed on your EWLM domain relative to time. The following are specific monitor views that are available: <ul style="list-style-type: none"> <li>▶ Goal achievement monitor</li> <li>▶ Processor utilization monitor</li> <li>▶ Transaction count monitor</li> <li>▶ Transaction rate monitor</li> </ul>

You can use any combination of these views to examine the performance of your EWLM domain. The views that you use to solve a problem are specific to your environment.

- ▶ For example, if a service class is missing its goal, the transaction class details reports that are specific to that service class can be useful. These reports specify the number of transactions associated with the transaction class and how they are performing.
- ▶ Also, if there is a problem with a particular application or server, the server topology or application topology can be useful. These topologies provide details about each application or server in the EWLM.

### 19.5.3 Monitoring scenario

Let's go through a scenario to see how EWLM can help diagnose a performance problem: To verify if the service classes are meeting their goals, you can use the Exceptions monitor. Figure 19-3 shows an example of the Exception monitor: If there are no entries in this section, all service classes are meeting the business response time goals you have set, otherwise you might want to start to investigate.



Select	Service class	PI	Importance	Performance	Goal
<input type="radio"/>	Response time service class	2.00	Highest	02.000 average	01.000 average
<input type="radio"/>	Percentile service class	0.00	Medium	0% in 45.000	60% in 45.000
<input type="radio"/>	EWLM Service Class	<None>	<None>	<None>	Discretionary

Page 1 of 1      Total: 3   Filtered: 3   Displayed: 3

Figure 19-3 Exception monitor

If the PI (Performance Index) is greater than 1, the service class goal is not being met. There can be numerous problems that can contribute to a goal not being met. For example, the processor utilization is too high, the number of transactions being processed is too high, the work requests are not classified to the correct service class, or the active service policy is not correct. You need to use a combination of detail reports, monitors, and topologies to determine why a service class goal is not being met. The specific monitor views that you use depend on the type of performance data that you obtain as you examine the problem.

In our sample, the performance data indicates that the Response time service class has a PI of 2.00. To determine why this service class goal is not being met, begin by viewing the details of the service class. In the details report, shown in Figure 19-3, examine the transaction statistics and identify which transaction classes are associated with this service class.

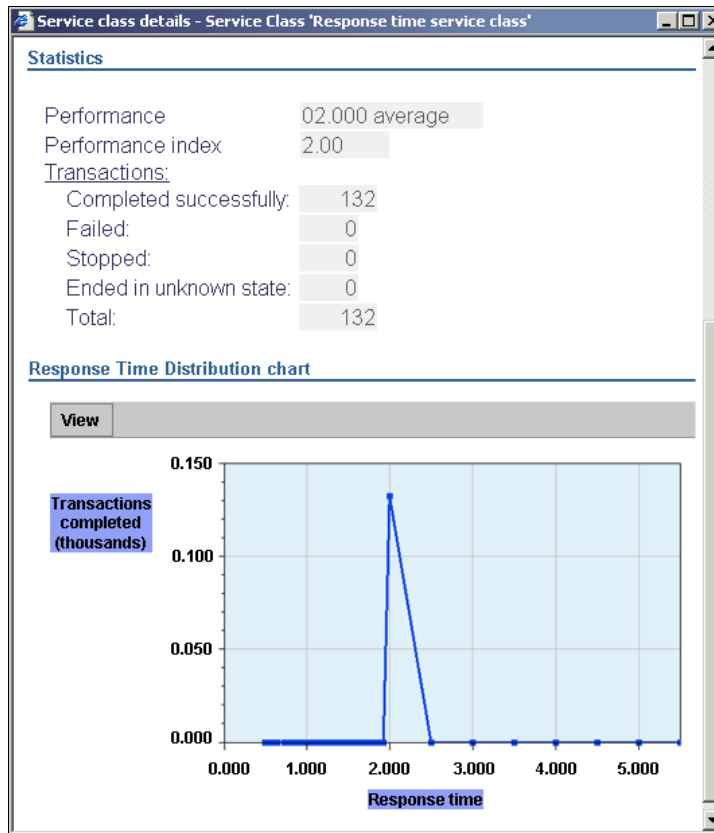


Figure 19-4 Service class details

Based on this data, you can determine that all of the transactions have completed successfully. In addition, you determine that only one transaction class is associated with this service class. This data does not indicate why the service class goal is not being met. The Response time distribution chart in the transaction class details report or service class details report helps you determine if any transactions have differing response times, and so, how many have a different response time and what is the response time? After viewing the transaction class details and service class details, you determine that all work requests take two seconds to complete. This data, however, still does not indicate the cause of the problem.

To investigate even further, view the server topology for the service class. The PI might be high due to problems with the managed server's processors that are processing the work requests. The server topology displays a high-level view of the servers in the EWLM domain. The server topology includes the managed servers and the ARM-instrumented applications which are processing work requests that are classified to the Response time service class.

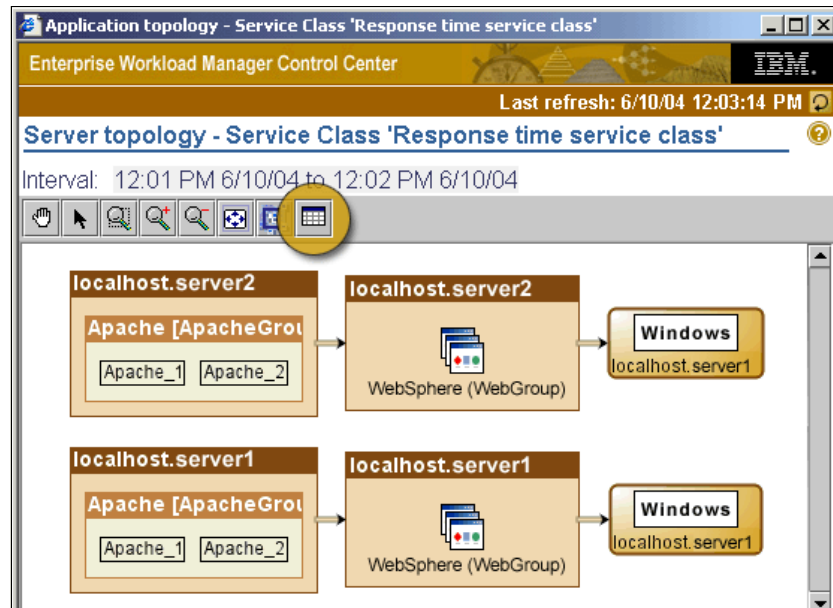


Figure 19-5 Server topology for the Response time service class

To obtain more specific data when you view a topology in the EWLM Control Center, you can click the table icon. The icon is highlighted in Figure 19-5, which shows the detailed performance statistics for this server topology.

Hop	Name	Type	Average response	Average active time	Processor using (%)
0	localhost.server1	Server	00.000	00.000	79%
0	Apache [ApacheGroup00]	Application	02.000	00.100	79%
0	Apache_1	Instance	02.000	00.100	79%
0	Apache_2	Instance	02.000	00.100	79%
0	localhost.server2	Server	00.000	00.000	56%
0	Apache [ApacheGroup00]	Application	02.000	00.100	57%
0	Apache_1	Instance	02.000	00.100	55%
0	Apache_2	Instance	02.000	00.100	55%
1	localhost.server1	Server	00.000	00.000	79%
1	WebSphere [WebGroup]	Application	01.900	00.400	79%
1	WebSphere_1	Instance	01.900	00.400	81%
1	WebSphere_2	Instance	01.900	00.400	81%
1	localhost.server2	Server	00.000	00.000	56%
1	WebSphere [WebGroup]	Application	01.900	00.400	55%
1	WebSphere_1	Instance	01.900	00.400	60%
1	WebSphere_2	Instance	01.900	00.400	60%
2	localhost.server1	Server	00.000	00.000	79%
2	DB2 [DBGroup]	Application	01.500	01.500	79%
2	DB2_1	Instance	01.500	01.500	79%
2	DB2_2	Instance	01.500	01.500	79%
2	localhost.server2	Server	00.000	00.000	56%
2	DB2 [DBGroup]	Application	01.500	01.500	57%
2	DB2_1	Instance	01.500	01.500	55%
2	DB2_2	Instance	01.500	01.500	55%

Figure 19-6 Server topology data

In this example topology data, you are particularly interested in the Average active time because work requests associated with this service class are completing, on average, in two seconds rather than one second. The Average active time specifies the amount of time, on average, that each server or application uses to process work requests. Use this data to determine if one of the server or applications is processing the work requests slower than expected. As shown in the server topology data, all database application instances in hop 2 have an average active time of 1.5 seconds. This is extremely high when compared to the goal of 2 seconds.

At this point, you should examine why the DB2 application appears to take longer to process than expected. View the Processor using % in the server topology to view the percentage of time the application was able to use the processor when requested. A high number indicates that the processor is available when needed. If the using percent is low, the service class is having trouble processing its work requests. In this example, localhost.server2 has a lower processor using percentage than localhost.server1.

To investigate the server (localhost.server2) with a lower processor using percentage further, look at the managed server details report to determine the Average processor utilization. This value specifies the total percentage of the all processors on the managed server that are being used. The following is example data from the managed server details for localhost.server2.



**Statistics:**

- ▶ Average processor utilization % 89
- ▶ Real memory 1,024
- ▶ Number of logical processors 4
- ▶ Page fault rate (faults per second) 8.0

**What are the recovery steps?**

Based on data that you obtained when investigating the problem, you determine that the main cause of the service class missing its goal is due to the processor utilization being too high. The server is potentially attempting to process too many work requests; therefore, the managed server cannot process the work requests in an efficient amount of time.

Here is a list of ways to recover:

- ▶ Move work requests to a different managed server for processing to reduce the processor utilization percent.
- ▶ Adjust the service class goal to a higher response time that takes into consideration a managed server with a high processor utilization percentage.
- ▶ Add more processor capacity to the managed server to reduce the total processor utilization of the managed server.
- ▶ Move high priority work requests to a different managed server. View the managed server details to determine what other service classes contain work requests that are being processed by the same server. Perhaps, you can move some work requests to a different service class.

Figure 19-7 is an example of the managed server's details report that identifies the service classes processed by the managed server.

Service classes processed in server	
Service class name:	EWLM Service Class
Processor utilization %	79.6
Processor delay %	15.9
Storage delay %	0.0
I/O delay %	0.0
Service class name:	Percentile service class
Processor utilization %	79.4
Processor delay %	15.9
Storage delay %	0.0
I/O delay %	0.0
Service class name:	Response time service class
Processor utilization %	79.6
Processor delay %	15.9
Storage delay %	0.0
I/O delay %	0.0

Figure 19-7 Managed server details report

View the Processor utilization field for each service class to determine how often work requests are able to process without being delayed. If the percent is low, the service class is not able to obtain processor resources when needed to complete its work requests. You should consider moving the service class to a different managed server for processing.

View the delay fields to determine if a large percent of a service class' work requests are being delayed. You should consider moving the service class to a different managed server for processing.

### 19.5.4 How to use EWLM to debug/diagnose a performance problem

In this section we want to discuss monitoring performance data using the EWLM Control Center focusing on how to use the reports and monitors to find a performance problem or bottleneck within the Web banking application topology.

The available monitoring and reporting functions are useful to determine whether or not there is a performance problem for a service class. They also provide a view of the logical tiers and the application environment and help to determine if applications or servers are experiencing any problems. The statistics include detailed resource usage and delay information, as well as correlation to specific operating system processes that support each middleware application.

### 19.5.5 Benefits and summary

Once fully deployed, EWLM is one element of the transformation of your IT infrastructure into an On Demand Business. It is an expansion of the concepts introduced by the z/OS Workload Manager over ten years ago, applying proven and mature technology to the multi-tier, heterogeneous environment. First and foremost, it helps your IT infrastructure “think” like you do, like a business. Starting from service level objectives established for your business applications, it learns the relationships between the servers, the application middleware instances and your business processes. Reporting information from a business perspective lets you rapidly understand when goals are not being achieved, where the bottlenecks are, and what the business impact is.

Detailed internal performance statistics are maintained by EWLM, explaining where time was spent and why, so you can quickly zoom into the area where investigation is required. It will also help you avoid investigating non-problem situations. The statistics maintained by EWLM form the basis for the introduction of goal-oriented autonomic management techniques, such as those from the z/OS platform. So, while you are looking at a potential problem, the environment is adapting itself to mitigate, or even eradicate the problem.

As was the case with the z/OS Workload Manager, the EWLM product introduces one level of management in the first release to get you started. Over time EWLM support will expand to internally automate management of server and application performance. Adaptive heuristic management is a necessary element of being “on demand”: avoiding problems, repairing problems, and helping you to raise server utilization to respectable levels

MyCom, Inc. now can enjoy the following benefits:

- ▶ End-to-end business transactions monitoring that help to understand the multi-tiered heterogeneous application and server topology and to understand the performance data for the applications.
- ▶ They can now understand the usage of the IT infrastructure and adjust the resources according to defined business policies to achieve desired business goals.
- ▶ By understanding the resources needed for the business applications, utilization of under-utilized computing systems is improved without requiring an investment in additional capacity, and the company’s systems “sense and respond” to disruption or threats before they occur.

## 19.6 Glimpse of the future

As technologies and products evolve, some of the items we might expect to see include having the same flexibility on all hardware platforms, regarding processor (including subprocessor), memory, and I/O allocation, and the number of independent operating systems hosted.

The workload manager that today runs within the operating system may find its way to the hypervisor to be able to balance resource across partitions, and eventually across different systems and architectures.

The monitoring capability introduced today by the workload manager can be seen as the initial step toward an self-managing autonomic computing environment where the performance goals set today for the monitoring arena will be extended to a management scope.

## 19.7 Summary

In this chapter we have explained how EWLM can provide a solution for monitoring business applications, driven by the installation business requirements. We have described the how the product works today and which feature are available right now that could be used to solve performance monitoring in complex enterprise environments. We also have presented an implementation scenario and an additional scenario to demonstrate better the value of the solution.



# Part 4

## Appendixes

This part of the book includes the following supplementary information:

- ▶ Appendix A, “Getting Started with the Virtualization Engine” on page 325
- ▶ Appendix B, “Standards overview” on page 333





# Getting Started with the Virtualization Engine

This chapter identifies the tools and services available inside IBM to learn more about the VE and to be able to evaluate, to prototype, to architect, and to implement the VE products.

Figure A-1 is a general representation of the usual tasks before implementing a new system application. Though every case is different, in activities, in duration, and in complexity, every new system application goes more or less through these high level multiple steps:

- ▶ The learning phase starts with presentations, sometimes with demonstrations, and then with predefined exercises to touch and feel the products.
- ▶ The evaluation phase usually includes a site assessment to understand how the site is well prepared to receive the new technology, a prototype to understand all the benefits of the new technology, inside or outside the customer site, and a cost evaluation to understand the cost benefits and the return of investments of the implementation.
- ▶ The implementation phase usually starts with a pilot to validate all the processes before a full implementation.

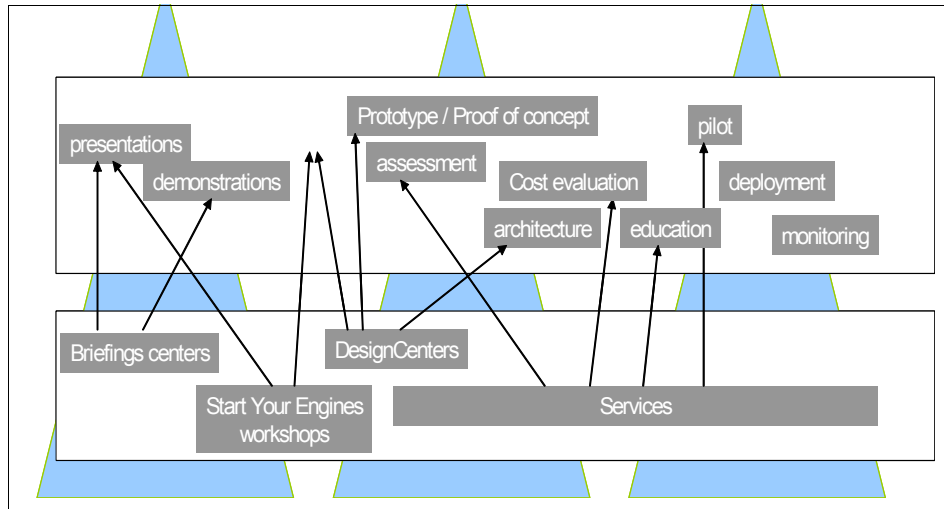


Figure A-1 Possible tasks before implementation

This appendix does not describe in details all these activities: some of them are very classical and very well known, some of them are short-lived by nature, such as the demonstrations which evolve with the availability of the concepts and features (some information about demonstrations available at the writing time in appendixX). This chapter focuses on specific activities less known and more specific to the VE environment such as the Start Your Engines workshops and the different types of assessment. A interactive tool is also available to help plan the ordering and installation processes.

Understand these offerings may be changed and customized according to specific customer needs and local implementation. The description of these offerings here is generic and wants to describe how to start a project that includes virtualization facilities. The VE being a new type of packaging embracing multiple technical layers, we think this is important to give the readers not only the concept and product descriptions but also a summary (and only a summary) of the activities available to help understand and implement these new components and how they work altogether. Specific offerings may exist for specific environment; do not hesitate to ask your IBM representative.

## A.1 “Start Your Engines” Workshops

To demonstrate the functions and the value of the VE components, a specific workshop, called “Start Your Engines”, has been developed, available in all IBM geographies.



Each 3 day workshop is a mix of lectures and hands-on practice, and is taught in different locations. Ask your IBM representative for the dates and locations.

The objective of the workshop is to provide an introduction to the IBM Virtualization Engine Suite, familiarize the participants with the strategy behind virtualization, and provide hands-on lab exercises so key aspects of the IBM Virtualization Engine Suite may be experienced and better understood. Several lab teams, every of each with its own multiple servers environment, are available for each workshop.

Figure A-2 describes a typical environment on which each student team can experiment the VE components.

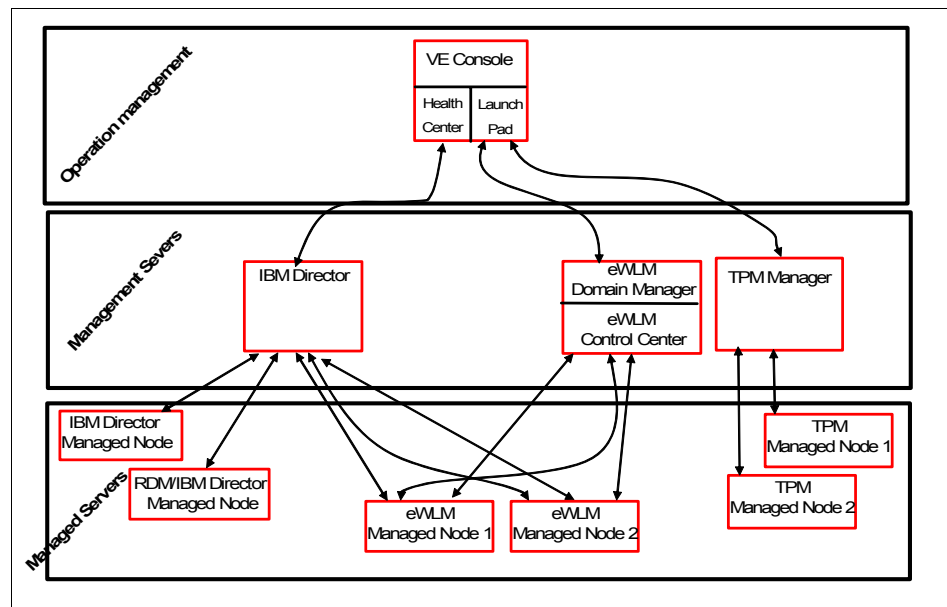


Figure A-2 Start Your Engines lab environment

At the time of writing, a typical workshop content is similar to the following list:

- Overview and strategy of IBM Virtualization Engine and the IBM Virtualization Engine Suite with a focus on the IBM Virtualization Engine Suite components
- IBM Virtualization Engine Console
- IBM Director Multiplatform
- IBM eWLM + Hands-on Lab
- IBM Tivoli Provisioning Manager
- IBM Grid
- IBM Storage Virtualization

Over time, the intent is to update the workshop by including the latest announced VE components and supported operating systems and platforms.

## **A.2 VE Consulting/Assessment study engagement overview**

Multiple services offerings covering the education, the consulting, the architecture, the implementation, and the management phases are available. This section covers the assessment offerings only.

You can get the latest details about the following offerings through the IBM local team.

### **A.2.1 Objectives**

The overall objective of a VE Consulting/Assessment project is to define the current state of a specific customer's IT server infrastructure, to describe realistic alternative future states and actions to reach these states, to evaluate the cost of the alternatives and to formulate a long term VE architecture and implementation strategy.

This is done by analyzing the information supplied by the customer about the current server infrastructure, by estimating its costs, and by developing alternative server solutions and business cases based on the new technologies.

The activity is planned to be of short duration with a tightly focused scope to minimize resource investment and quickly produce the key elements of the architecture.

The consulting and assessment offerings are designed to assist the customer with the identification of VE technologies and solutions to meet specific business and technology pain points.

## A.2.2 Offerings

Figure A-3 describes the different types of assessment offerings.

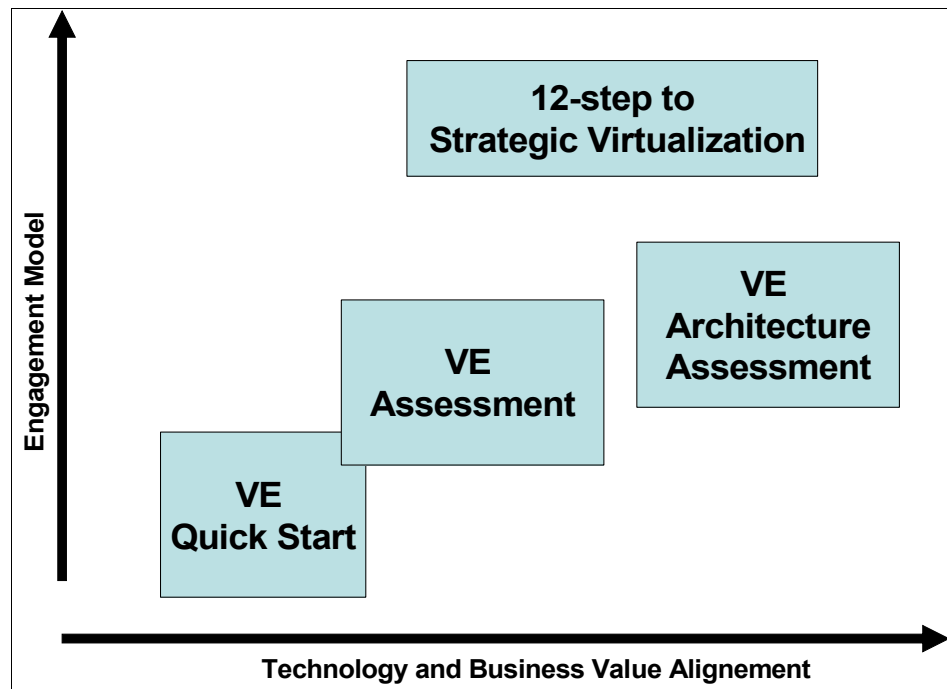


Figure A-3 Virtualization Engine Assessment offerings

### A.2.2.1 The VE Quick Start Assessment

The VE Quick Start consulting offering is a 2 day duration study.

It is very effective in situations where a targeted look at a particular subset of requirements is appropriate.

The study identifies high level business strategy and knowledge of customer pain-points; it identifies appropriate customer solutions and VE technologies; it provides quantified benefits in terms of TCO and ROI through the reference architectures.

The findings of the project enable IT Executives to balance alternative VE based technical options and make sustainable server platform and software component decisions.

#### **A.2.2.2 The VE Assessment**

The VE Assessment offering is a 5 to 10 day duration study.

It is appropriate for analyzing the customer current state and future business and IT strategy, the current implementation and the whole server infrastructure.

This is especially valuable where it is appropriate to envisage a technology paradigm shift.

The study identifies business, strategy and knowledge of customer pain-points; it translates that knowledge into VE technology solutions; it identifies appropriate customer specific product adoption strategy and provides quantified benefits in terms of TCO and ROI.

#### **A.2.2.3 The VE Architecture Assessment**

The VE Architecture Assessment offering is a 2 to 10 days duration study.

This study is appropriate to insure that the introduction of the VE components will fit with the business plan, the technologies capabilities and future customer strategies.

The study evaluates the hardware and software structures of the IT system solution, including components that will be reused, developed or purchased. It ensures the new system will support the users and the required business functionality.

#### **A.2.2.4 The VE Consulting Assessment**

The VE Consulting Assessment (sometimes known as the 12-step to Strategic Virtualization offering) is a 2 days to several weeks duration study.

This offering is designed to assist customers bridging the gap from the current state of their current business and IT environment towards their strategic business and technology objectives. It leverages several requisite and measurable solutions and technologies through infrastructure simplification solutions such as IT optimization, server consolidation, and virtualization technologies.

### **A.2.3 Methodology**

The study is typically delivered as a comprehensive project by investigating an enterprise's total IT server infrastructure and storage hardware, software and service delivery staff.

Though it will depend on the type of assessment, the project usually is comprised of different phases:

- ▶ A first phase evaluates customer's IT state to understand the current technology implementation, the requirements, the challenges, and the business strategy. The objective is to give to all participants a common view and understanding of the environment and of the engagement.
- ▶ Another phase reviews the current state of virtualization technologies, capabilities and benefits; implementation procedures and challenges, integration points of available VE technologies and systems topologies are reviewed in details as well.
- ▶ The next phase purpose is to review the integration of particular and specific VE technologies, capabilities and benefits, looking specifically at the architectural and implementation procedures as it pertains to the specific environment. In this phase, time and cost implementation, service level benefits, impact of the integration of future processes and application are estimated.
- ▶ The last phase reviews the VE global solution design and validates the technical design of the VE components in terms of the global solution to ensure that the plan will be successfully implemented. It includes a technical detailed review and validation. It includes a review to confirm that the solution will meet customer expectations and that associated risks can be adequately addressed.

For each phase, interviews and workshops will be proposed. The purpose of these activities is to provide the engagement team with the information needed to understand the environment and the challenges and to provide the necessary education to be sure the VE components functions are well understood.

Reports are distributed periodically during the engagement, and of course at the end of the engagement. They consist of observations, findings, and propositions in the area, usually, of technology, finances, and strategy. Recommendations are prioritized and next steps can be planned.

## A.3 The Planning Advisor

Before implementing the IBM Virtualization Engine, there are important ordering, planning, and deployment considerations to be made. The Virtualization Engine planning advisor is an interactive tool that produces customized planning output based on customers needs.

To access the tool, contact your IBM representative.

Consider important planning, ordering, and topology issues before deploying the IBM Virtualization Engine. The Virtualization Engine planning advisor interview will ask a series of questions about the environment and how you plan to use the Virtualization Engine.

Based on the input, the planning advisor generates recommendations to help you:

- ▶ **Order:** The planning advisor interview helps you decide which systems services to include in your Virtualization Engine enterprise
- ▶ **Deploy:** Once you determine which systems services meet your business needs, the planning advisor recommends how to deploy the components of the Virtualization Engine across your enterprise topology.
- ▶ **Install:** Finally, the planning advisor generates customized documentation to help you prepare the environment, and to help you install and configure the Virtualization Engine.

The planning advisor interview provides you with information about the capabilities and requirements of the Virtualization Engine systems services to help you determine which pieces of the Virtualization Engine you would like to order and install.

Then it describes a topology that represents a logical representation of where the systems service components of your Virtualization Engine deployment should be installed in your enterprise.

Then it provides charts to guide you through the installation and configuration of each of the systems services you selected.



# B

## Standards overview

An on demand computing environment involves diverse systems working together and connecting to devices and applications across platforms, organizations, and even geographic borders. This environment helps to enable a business to respond quickly to changes in markets, technologies, and the needs of their customers. Businesses must be able to rapidly provide new capabilities to their systems without completely discarding and replacing those applications and systems.

The only way all of these components, applications and systems can work together is with open industry standards. With open standards, businesses and providers can ensure that their products and systems will work and communicate with other systems.

Standards enable an On Demand Business to create the environment and applications needed. IBM is actively involved in developing many of these standards, as are many other companies worldwide.

## B.1 Open source

Open standards are necessary to open-source projects. Open-source projects frequently provide implementations of key standards that serve as references. The Apache Web server is one example. Because the Apache Web server is so commonly used, every browser must work with its implementation of the Hypertext Transfer Protocol (HTTP) standard. This creates a market pressure that prevents vendors from introducing incompatible, proprietary implementations of HTTP.

IBM contributes to open-source projects, and actively supports the open-source community. IBM contributed an Extensible Markup Language (XML) parser to the Apache Xerces project, and an XML Stylesheet Language Transformation (XSLT) engine to the Xalan project. In addition, IBM created the Eclipse project, an effort to create an open-source integrated development environment (IDE). In turn, many open-source tools are being incorporated into IBM development tools.

## B.2 Standards organizations

There are many standards organizations contributing to the key standards for on demand computing. The following are some of the key standards organizations that one should be aware of when planning for an On Demand Operating Environment.

- ▶ Internet Engineering Task Force (IETF)
- ▶ World Wide Web Consortium (W3C)
- ▶ Java Community Process (JCP)
- ▶ Organization for the Advancement of Structured Information Standards (OASIS)
- ▶ Web Services Interoperability Organization (WS-I)
- ▶ Distributed Management Task Force (DMTF)
- ▶ Global Grid Forum (GGF)
- ▶ Object Management Group (OMG)

### B.2.1 IETF - Internet Engineering Task Force

The Internet Engineering Task Force (IETF) creates standards for the operation of the Internet infrastructure. The IETF was formed in 1986 and has evolved into an active standards organization involving thousands of people from academia, research, and industry. The IETF has no formal membership. Anyone may participate in mailing lists or attend meetings. Participants are organized into an ever-changing collection of Working Groups, which are further organized into Areas. While IETF Working Groups can be created in any area based on the interests of the participants, in practice, the IETF concentrated on the



transmission of Internet Protocol (IP) packets and the information required to secure, route, and manage the communications.

Because on demand computing assumes computer networking as a base capability, almost all IETF standards have an impact. A unique requirement of on demand computing, though, is highly scalable and dynamic networking. The IETF is leading the transition of the Internet infrastructure to a new base protocol, known as IPv6, which will dramatically increase the number of Internet addresses available and simplify address management. In addition, the IETF continues to evolve security and routing protocols that enable dynamic creation of secure networks.

## **B.2.2 W3C - World Wide Web Consortium**

The World Wide Web Consortium (W3C) creates specifications for Web technologies. The mission of W3C is to lead the Web to its full potential. It does this by developing recommendations, guidelines, software and tools that create a forum for information, commerce, inspiration, independent thought, and collective understanding.

Tim Berners-Lee, who is widely credited as being the architect of the World Wide Web, founded the W3C to further the growth of the Internet and ensure its interoperability. The W3C is a consortium of companies working together to develop Web technologies. HTML, Extensible Hypertext Markup Language (XHTML), XML and XML Schema are just a few examples of W3C Recommendations.

The W3C organizes the work on the development of Web technologies into Activities. The structures of these Activities vary, but each Activity usually includes a Working Group, an Interest Group, and a Coordination Group. Within their respective Activities, these groups produce Recommendations and other technical reports as well as sample code.

W3C activities of interest include the following:

- ▶ XML family of standards (the XML, XML Base, XML Query, XML Schema, XPath, and XSLT standards are of particular interest)
- ▶ Simple Object Access Protocol (SOAP)
- ▶ Web Services Description Language (WSDL)

## **B.2.3 JCP - Java Community Process**

The Java Community Process (JCP) organization, created by Sun Microsystems, was formed to create and maintain Java technical specifications.

The companies and Java developers who make up the JCP provide specifications, reference implementations, and technology compatibility kits to guide the evolution of Java technology. The open organization works with member and nonmember input. Anyone can join the organization, but membership is not required in order to contribute.

The JCP works to ensure the stability and cross-platform compatibility of Java. It also works to expand the platforms specification portfolio to address emerging technologies.

## **B.2.4 OASIS - Organization for the Advancement of Structured Information Standards**

In 1993, a consortium of vendors and users formed SGML Open to develop guidelines for interoperability among products that support the Standard Generalized Markup Language (SGML). By 1998, the scope of the work had expanded to include XML and other related standards, and the name was change to the Organization for the Advancement of Structured Information Standards (OASIS).

Today, OASIS is a not-for-profit, global consortium driving the development, convergence, and adoption of e-business standards. OASIS develops structured information standards for security, Web Services, XML conformance, business transactions, electronic publishing, topic maps, and interoperability within and between marketplaces. OASIS members set the technical agenda, following a process designed to achieve industry consensus and unite disparate efforts.

Key OASIS specifications of interest to the IBM community include:

- ▶ UDDI
- ▶ Web Services Security (WS-Security)
- ▶ Business Process Execution Language (BPEL4WS)

## **B.2.5 WS-I - Web Services Interoperability Organization**

The Web Services Interoperability Organization (WS-I) promotes interoperability between Web Services across platforms, applications, and programming languages. The organization includes diverse group of software vendors, enterprise customers, and others interested in Web Services to aid the development of interoperable Web Services with guidance, recommended practices, and resources. The WS-I provides resources to help developers create Web Services that are interoperable and compliant with WS-I guidelines and industry standards.

The WS-I Basic Profile 1.0 specification describes ways in which diverse Web Services specifications can work together to create interoperable Web Services. The WS-I is also working on a profile to cover the implications and workings of the OASIS standard, WS-Security.

## **B.2.6 DMTF - Distributed Management Task Force**

The Distributed Management Task Force, Inc. (DMTF) promotes the development and adoption of interoperable management standards for enterprise and Internet environments. They developed the Common Information Model (CIM) standard, which describes a platform-independent method for exchanging management information. The standard helps to simplify integration and reduce costs for management systems by enabling an end-to-end multi-vendor interoperability. By implementing CIM, vendors and standards groups make possible more integrated and cost-effective management systems.

DMTF standards with which you should be familiar for on demand computing are:

- ▶ CIM
- ▶ Web Based Enterprise Management (WBEM)

## **B.2.7 GGF - Global Grid Forum**

The Global Grid Forum (GGF) promotes and supports grid technologies and applications. They create specifications, user experience documents and implementation guidelines to help organizations developing, deploying and implementing grid technologies.

In addition, they promote the development of a broad-based Integrated Grid Architecture to support emerging grid communities. Such architecture can aid the grid agenda by spreading necessary basic services and encouraging the use of shared middleware code for applications with common requirements.

GGF recommendations with which you should be familiar for on demand computing include:

- ▶ Open Grid Services Infrastructure (OGSI), a base set of distributed computing operations to support dynamic middleware
- ▶ Open Grid Services Architecture (OGSA), a model of a computing system as a set of distributed computing patterns realized as applications and extensions of Web Services
- ▶ Distributed Resource Management Application API (DRMAA), an application programming interface specification for the submission and control of jobs to one or more Distributed Resource Management (DRM) systems

## B.2.8 OMG - Object Management Group

The Object Management Group (OMG) is a nonprofit consortium whose purpose is to promote object-oriented technology and the standardization of that technology. The OMG was formed to help reduce the complexity, lower the costs, and accelerate the introduction of new software applications. Some of OMG's major successes include the Common Object Request Broker Architecture (CORBA), and the Unified Modeling Language (UML). One of OMG's current efforts is establishing the standards for Model-Driven Architecture (MDA).

OMG specifications with which you should be familiar for on demand computing include:

- ▶ MDA
- ▶ UML

Although W3, OASIS, IETF and OMG are key standards-setting bodies for our future grid services world, it is important for developers to follow the interoperability standards set by the WS-I. Web Services support and make possible key elements of the emerging grid services.

## B.3 Key standards

At least 158 standards are significant to the on demand strategy. These standards apply to any of 22 different categories, including messaging, security, management, Java, and discovery categories, to name just a few.

There are also "vertical" standards that support the on demand strategy. Vertical standards refer to business standards or regulations that developers must follow when developing software for particular sectors or industries. Examples of these vertical standards include RosettaNet for electronics and ACORD for insurance.

The following are some of the key standards that apply to an On Demand Operating Environment:

- ▶ XML standards, including XML Schema and XSLT
- ▶ SOAP
- ▶ WSDL
- ▶ UDDI
- ▶ WS-I Basic Profile
- ▶ WS-Security
- ▶ OGSA
- ▶ OGSF
- ▶ UML

- ▶ MDA
- ▶ WBEM
- ▶ CIM, CIM-XML, CIM-SOAP

Of these, SOAP, WSDL, UDDI, WS-I Basic Profile, and WS-Security are basic Web Services standards. Figure B-1 shows how these and other Web Services standards relate to one another.

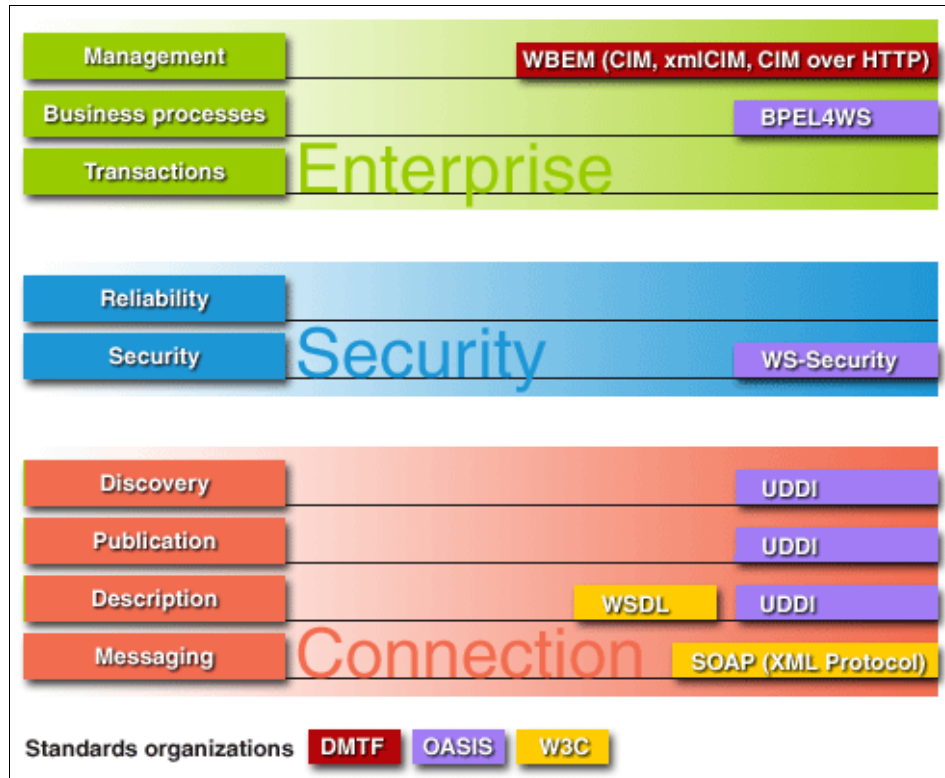


Figure B-1 Key standards

### B.3.1 XML standards

XML is a family of technologies for defining and processing application-specific markup languages that describe data and documents. Specifically, XML is a language for creating and using structured information. XML is based on, and is a subset of, SGML.

In simple terms, XML is simply a standardized format for the representation of data documents. It was developed by an XML Working Group formed under the auspices of the W3C in 1996 and provides the foundation for many of the open

standards of today. This is particularly true of those standards related to the interoperability of applications and components, such as WSDL, since XML defines a simple base structure for the representation of data.

Resources on the details of XML syntax and related technologies are numerous.

Some of the inherent benefits of XML include:

- ▶ It is an easy-to-use, open standard for data description and as such, forms a convenient common ground between heterogeneous applications and components.
- ▶ Its element-oriented structure means that XML is indeed easily extensible. A common problem with proprietary file formats (such as fixed-width record files) is that they are often only able to withstand a finite amount of extension (lack of space in the record, for example). The tag structure of XML makes the addition of new tags and attributes straightforward.
- ▶ XML documents are, generally speaking, easier for humans to read and understand (and therefore, debug or analyze) than comma-separated or hash-delimited files. For example, compare the following data formats, which relate to the same piece of data:

In a hash-delimited format:

```
1#martin#gale#d0168
```

In XML:

```
<employee id="1">  
<name><forename>martin</forename><surname>gale</surname></name>  
<office>d0168</office>  
</employee>
```

- ▶ XML defines languages for describing the structure of a particular XML document in order for it to be valid in terms of its application. XML standards describe the syntactical constructs for the base markup of a document. The validation uses a Document Type Definition (DTD) document or, more recently, an XML Schema document, both of which describe the validation rules for the data. DTDs and XML Schemas are referenced from within a given XML document using a Uniform Resource Locator (URL). This allows the document to be validated regardless of the platform on which it is processed.
- ▶ A variety of freely available, open-source XML parsers for various programming languages make integrating structured data described in XML into an application straightforward. Likewise, the availability of XSLT processors means that the translation of XML from one format into another is a portable and straightforward process.

### ***XML Schema***

XML Schema is a key XML-related technology in the on demand world. XML Schemas express shared vocabularies and allow machines to carry out rules made by people. They provide a means for defining the structure, content and semantics of XML documents, and allow the entities within the XML data to be bound to user-defined type information. Schema is a modernization of the XML DTD principle that is in itself described in XML, as opposed to DTD's proprietary format.

As an XML vocabulary itself, Schema carries with it all the benefits of XML, particularly in respect to portability. In this way, XML Schema documents are used in conjunction with WSDL in Web Services to describe not only the Web Service provided, but to define the data types consumed by that service. Similarly, XML-based standards are now appearing defined in XML Schema. In these cases, the schema provides a single standard format artifact that describes the vocabulary in question, and is a valid run-time asset with which to validate the documents using it.

### ***XSLT***

Other key standards conveyed in XML are XSLT. In general terms, XSLT stylesheets describe the mapping of an XML document from one XML format to another. XSLT stylesheets can transform a document containing data into output markup in which the data is contained within formatting constructs (such as XHTML). In addition, where a key standard is expressed as an XML vocabulary, you can also use XSLT itself to generate new controlling documents, including XSLT stylesheets and XML Schema definitions. In many ways, XSLT is useful as a flexible, portable, and relatively easy to use markup.

## **B.3.2 SOAP - Simple Object Access Protocol**

SOAP is an XML-based protocol for applications to send and receive messages over the Internet. It is a W3C Recommendation. SOAP defines an “envelope” that allows clients and service providers to communicate and exchange XML-formatted data, regardless of platform or programming language. The specification defines the XML formatting for the messages, a method for encoding the data as XML and a binding to HTTP as the transmission method.

The specification allows for using other encoding methods, but this is discouraged because it would limit the potential users and potentially fragment the SOAP user base. The SOAP specification also allows other bindings, such as WebSphere MQ.

### **B.3.3 WSDL - Web Services Description Language**

WSDL is an XML format for describing the interfaces of a Web Service. The details described include the protocols and port numbers used, available operations, and message formats. WSDL is being defined by the W3C. In basic terms, the language is used to describe the capabilities of a Web Service (for example, the operations that can be performed). The details described include the protocols and port numbers used, message formats and possible exceptions.

IBM and Microsoft jointly developed WSDL. Note that the UDDI standard uses WSDL.

### **B.3.4 UDDI - Universal Description, Discovery, and Integration**

UDDI is a business-registry specification used to support the description and discovery of Web Services providers, including businesses and organizations, the Web Services offered by those providers, and the interfaces available to access those services. UDDI is an OASIS specification for indexing Web Services so that users can locate and use them.

UDDI makes use of several standards, including SOAP, XML Schema, and WSDL. It provides a mechanism for clients to find other Web Services. Entries in a UDDI registry include information on the business offering a Web Service, the capabilities of the services offered, and technical details on how to invoke and use the service.

### **B.3.5 WS-I Basic Profile 1.0a**

The WS-I Basic Profile Version 1.0a was released August 8, 2003. The Basic Profile describes a manner in which four key Web Services specifications can be implemented to achieve a consistent measure of interoperability. Those four specifications are:

- ▶ XML Schema 1.0
- ▶ SOAP 1.1
- ▶ WSDL 1.1
- ▶ UDDI 2.0

Other standards organizations manage these specifications. The Basic Profile does not add to the specifications. It seeks to show how they can work together.



Among other issues, the Basic Profile addresses the major interoperability concerns, provides a way of testing, clarifies the requirements in specifications (such as avoiding optional components as sources of possible confusion), and makes strong recommendations regarding multiple possible implementation mechanisms that may be found in the specifications themselves.

### **B.3.6 WS-Security - Web Services Security**

Web Services Security (WS-Security) is a proposed specification for enhancing the security of SOAP messaging. Developed jointly by IBM, Microsoft, and Verisign, WS-Security has been submitted to OASIS.

WS-Security provides a basic mechanism for linking security tokens to SOAP messages. Designed to support multiple types of security tokens, it does not specify the use of any particular one. WS-Security describes how the binary tokens should be encoded. By itself, WS-Security does not ensure security, but it is designed to make use of other Web Services extensions and higher level application-specific protocols.

### **B.3.7 OGSA - Open Grid Services Architecture**

The OGSA is a model of a computing system as a set of distributed computing patterns realized as applications and extensions of Web Services. IBM supports the development of OGSA for the management of a virtualized set of resources. The Global Grid Forum manages the OGSA effort.

OGSA defines the elements necessary to build and run a platform for distributed system integration. These elements include:

- ▶ The scope of the services needed to support scientific and business applications
- ▶ The core set of services necessary for grid systems and applications
- ▶ The functions needed for the core services and the relationships between them

OGSA integrates key grid technologies with Web Services mechanisms to create a distributed system framework based on the Open Grid Services Infrastructure (OGSI). Unlike OGSI, OGSA addresses the creation, management and destruction of grid services.

### B.3.8 OGSi - Open Grid Services Infrastructure

OGSi defines methods for the creation, management, and exchange of information between grid services. A grid service is defined as a Web Service that conforms to a set of conventions, expressed as WSDL interfaces, extensions, and behaviors. The elements address purposes such as lifetime management, discovery of characteristics, and notification. Grid services provide a method for managing the distributed and possibly prolonged state that is often required by complex distributed applications. OGSi also specifies methods for the creation and discovery of grid services with standard factory and registration interfaces.

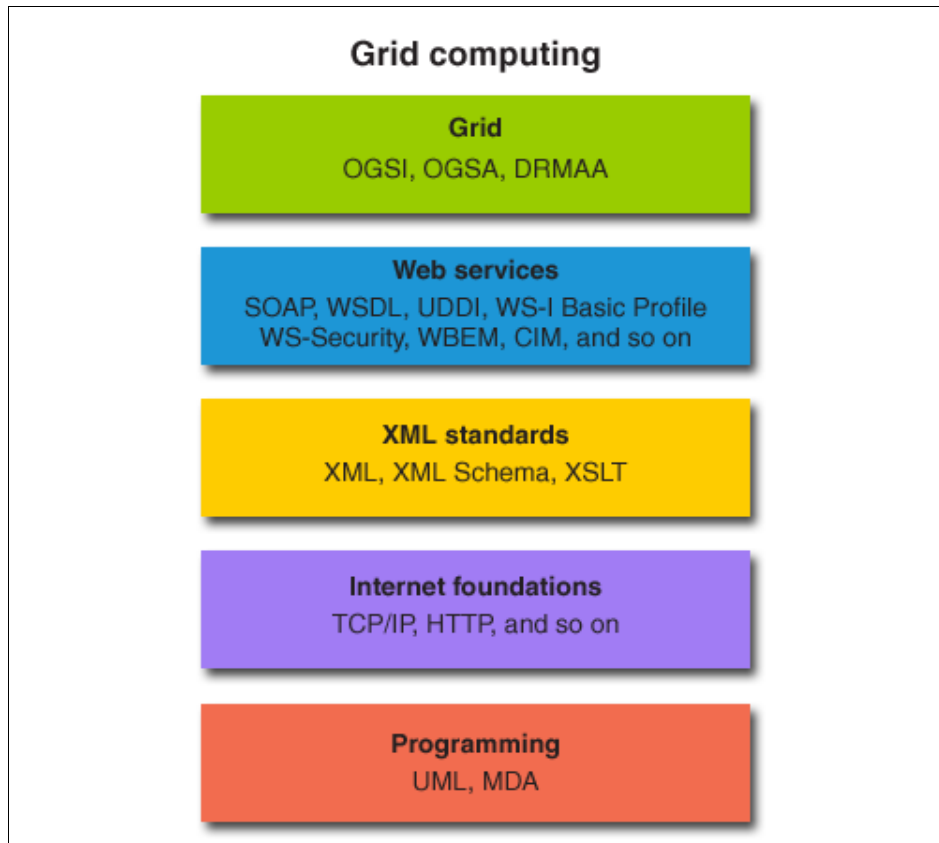
According to the current draft of the specification, the OGSi component model extends WSDL and XML Schema to incorporate:

- ▶ Stateful Web Services
- ▶ Inheritance of Web Service interfaces
- ▶ Asynchronous notification of state change
- ▶ References to instances of services
- ▶ Collections of service instances
- ▶ Service state data that augments the constraint capabilities of XML Schema definition

Note that OGSi defines the basic elements and mechanisms that OGSA uses to create a grid services platform.

DRMAA is a related specification being developed by the GGF for the submission and control of jobs to one or more DRM systems.

Figure B-2 illustrates how the grid specifications relate to other standards. Grid computing makes extensive use of Web Services. In turn, Web Services rely on XML. None of it would be practical without the underpinnings of the Internet and the programming that creates the applications involved.



*Figure B-2 Grid specifications and related standards*

### **B.3.9 UML - Unified Modeling Language**

UML is a standard notation used to visually design and model applications and systems. UML is a language and not a methodology, and as such, it is independent of programming languages or platforms. It can be used with any programming language to create design plans and illustrate system structures. UML diagrams and documents essentially serve as blueprints. Although UML is used extensively in application development, it can also be used for business modeling and other non-software types of systems. UML is an OMG specification and forms the basis of another OMG specification, MDA.

### **B.3.10 MDA - Model Driven Architecture**

MDA is an OMG specification based on UML. Making extensive use of UML models, MDA offers a way of creating specifications and developing applications separated from the underlying technologies of the platforms used. Leveraging UML, developers using MDA can design interfaces and relationships between applications independent of platforms and programming languages. The applications designed this way can be created on a range of platforms, open or proprietary. When new technology is developed, the modeling does not need to be repeated.

The Eclipse project is an example of an industry open model framework that is MDA compliant.

Using MDA, you start with a Platform Independent Model (PIM). This is a model of functionality and behaviors without details on the technical implementation. MDA-compliant tools then map the PIM to a Platform-Specific Model (PSM), or more likely, to multiple PSMs. This partially automatic process is accomplished using OMG-standardized mappings. The resulting PSMs are also UML models.

The MDA tools then use the PSM models to generate actual code, including interfaces, configuration files, and more. Depending on the complexity of the model or application, the tools generate all or most of the code needed. At this point, the code can be fine-tuned before the application is deployed.

### **B.3.11 CIM - Common Information Model**

CIM is a DMTF standard for expressing data about systems, applications, networks, and devices. CIM allows various management applications to access the data and control the devices or systems regardless of the platforms involved, making interoperability easier to achieve.

The CIM provides object classes, properties, methods, and associations common to the use of management applications in the form of management schema. These are organized into three layers:

- ▶ A core model addresses elements that span all areas of management.
- ▶ Common models address elements found in specific management areas, such as systems, applications, networks, and devices, independent of technologies or implementations.
- ▶ Extension schemas address the needs of specific technologies (specific operating systems or platforms).

CIM is independent of the method used for implementation.

The Web-Based Enterprise Management (WBEM) Initiative is an effort by the DMTF to design standards for the management of computing environments. The goal is to provide the standards that the industry can use to create integrated, standards-based tools that make use of Web technologies. At the core of this initiative is the CIM.

Related standards include:

- ▶ xmlCIM specifies a way in which CIM elements and messages can be expressed in XML.
- ▶ CIM Operations over HTTP specifies how to map CIM operations onto HTTP to achieve an open, standardized interoperability.

### **B.3.12 Open Group**

The Open Group is an international vendor and technology-neutral consortium that is committed to delivering greater business efficiency by bringing together buyers and suppliers of information technology to lower the time, cost and risk associated with integrating new technology across the enterprise.

With its proven certification methodology and conformance testing expertise, the Open Group is an international facilitator in delivering the interoperability that single organizations require to help ensure independence.

The flexible structure of membership of The Open Group allows for almost any size organization to join and influence the future of the IT world. The introduction of membership for individuals is currently being considered.

The Open Group offers the following services:

- ▶ Certification
- ▶ Testing
- ▶ Forums
- ▶ Initiatives
- ▶ The Open Brand
- ▶ Quarterly Conferences
- ▶ Regional Chapters
- ▶ Advanced Research
- ▶ Consortia Services



# Related publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this redbook.

## IBM Redbooks

For information on ordering these publications, see “How to get IBM Redbooks” on page 354. Note that some of the documents referenced here may be available in softcopy only.

- ▶ *VMware ESX Server:Scale Up or Scale Out?, REDP-3953*
- ▶ *Server Consolidation with VMware ESX Server, REDP-3939*
- ▶ *Volume Virtualization with the IBM SVC Software on the CISCO MDS 9000O, TIPS0392*
- ▶ *Server Consolidation: A Comparison of Workload Management and Partitioning, TIPS0426*
- ▶ *Introduction to pSeries Provisioning, SG24-6389*
- ▶ *Cluster Systems Management Cookbook for pSeries, SG24-6859*
- ▶ *Advanced Virtualization Features on p5 Servers, TIPS0477*
- ▶ *zSeries Application Assist Processor (zAAP) Implementation, SG24-6386*
- ▶ *Achieving the Highest Levels of Parallel Sysplex Availability, SG24-6061*
- ▶ *OS/390 Workload Manager Implementation and Exploitation, SG24-5326*
- ▶ *Introduction to pSeries Provisioning, SG24-6389*
- ▶ *Advanced POWER Virtualization on IBM eServer p5 Servers:Introduction and basic Configuration, SG24-7940*
- ▶ *IBM Enterprise Workload Management, SG24-6350*
- ▶ *Grid Computing with the IBM Grid Toolbox, SG24-6332*
- ▶ *Federated Identity Management with IBM Tivoli Security Solutions, SG24-6394*
- ▶ *On demand Operating Environment: Creating Business Flexibility, SG24-6633*
- ▶ *Tivoli Storage Manager Version 3.7: Technical Guide, SG24-5477-00*

- ▶ *IBM Tivoli Monitoring for Databases Database Management Made Simple*, SG24-6613-00
- ▶ *IBM Tivoli Intrusion Manager Version 3.7.1 Rapid Deployment Guide*, REDP-3695-00
- ▶ *Tivoli SecureWay® Risk Manager Correlating Enterprise Risk Management*, SG24-6021-00
- ▶ *e-business Risk Management with Tivoli Risk Manager*, SG24-6036-00
- ▶ *Introducing IBM Tivoli Service Level Advisor*, SG24-6611-00
- ▶ *Up and Running with WebSphere Site Analyzer*, SG24-6169-00
- ▶ *Introducing IBM Tivoli License Manager*, SG24-6888-00
- ▶ *Automated Distribution and Self-Healing with IBM Tivoli Configuration Manager Version 4.2*, SG24-6620-00
- ▶ *Unveil Your e-business Transaction Performance with IBM TMTP 5.1*, SG24-6912-00
- ▶ *IBM Tivoli Access Manager for e-business*, REDP-3677-00
- ▶ *Enterprise Security Architecture using IBM Tivoli Security Solutions*, SG24-6014-00
- ▶ *Enterprise Business Portals with IBM Tivoli Access Manager*, SG24-6556-00
- ▶ *Enterprise Business Portals II with IBM Tivoli Access Manager*, SG24-6885-00
- ▶ *e-business On Demand Operating Environment*, REDP-3673-00
- ▶ *Tivoli Storage Manager Version 3.7: Technical Guide*, SG24-5477-00
- ▶ *z/OS Intelligent Resource Director*, SG24-5952-00
- ▶ *Implementing Systems Management Solutions using IBM Director*, SG24-6188-01
- ▶ *IBM Tivoli Monitoring for Databases Database Management Made Simple*, SG24-6613-00
- ▶ *IBM Web Infrastructure Orchestration*, SG24-7003-00
- ▶ *Introducing IBM Tivoli On Demand Provisioning Solutions*, SG24-8888-00

## Other publications

These publications are also relevant as further information sources:

- ▶ *IBM Tivoli Intelligent Orchestrator and IBM Tivoli Provisioning Manager Overview Guide Version 1.1.0*, SC32-1419-00



- ▶ *IBM Tivoli Access Manager for WebSphere Application Server User Guide Version 4.1*, SC32-1136-00
- ▶ *IBM Tivoli Access Manager for WebLogic Server User's Guide Version 4.1*, SC32-1137-01
- ▶ *Tivoli Identity Manager Policy and Organization Administration Guide Version 4.5.0*, SC32-1149-01

## Online resources

These Web sites and URLs are also relevant as further information sources:

- ▶ About Virtualization Solutions:  
<http://www-1.ibm.com/servers/eserver/about/virtualization/>
- ▶ IBM Virtualization Engine:  
[http://www-1.ibm.com/servers/eserver/about/virtualization/getting\\_started.html](http://www-1.ibm.com/servers/eserver/about/virtualization/getting_started.html)
- ▶ IBM TotalStorage Virtualization Solutions:  
<http://www-1.ibm.com/servers/storage/software/virtualization/library.html>
- ▶ Virtualization Engine Planning Advisor:  
[http://publib.boulder.ibm.com/infocenter/eserver/vir1/en\\_US/index.htm?info/veicinfo/eicarplangeneral.htm](http://publib.boulder.ibm.com/infocenter/eserver/vir1/en_US/index.htm?info/veicinfo/eicarplangeneral.htm)
- ▶ zSeries Virtualization:  
<http://www-1.ibm.com/servers/eserver/zseries/virtualization/features.html>
- ▶ iSeries Virtualization:  
<http://www-1.ibm.com/servers/eserver/about/virtualization/systems/iseries.html>
- ▶ pSeries Virtualization:  
<http://www-1.ibm.com/servers/eserver/about/virtualization/systems/pseries.html>
- ▶ xSeries and Blade Center Virtualization:  
<http://www-1.ibm.com/servers/eserver/about/virtualization/systems/xseriesbladecenter.html>
- ▶ Infrastructure Simplification:  
<http://www-306.ibm.com/software/tivoli/solutions/resource/>
- ▶ IBM On demand Web pages:  
<http://www.ibm.com/ondemand>
- ▶ How IBM Infrastructure Orchestration works:  
<http://www.ibm.com/news/us/2003/10/081b.html>

- ▶ IBM Tivoli Provisioning Manager:  
<http://www-3.ibm.com/software/tivoli/products/prov-mgr/>
- ▶ IBM Tivoli Intelligent ThinkDynamic Orchestrator:  
<http://www-3.ibm.com/software/tivoli/products/intell-orch/>
- ▶ IBM Tivoli Monitoring:  
<http://www-3.ibm.com/software/tivoli/products/monitor/>
- ▶ IBM Tivoli Risk Manager:  
<http://www-3.ibm.com/software/tivoli/products/risk-mgr/index.html>
- ▶ IBM Tivoli Intrusion Manager:  
<http://www-3.ibm.com/software/tivoli/products/intrusion-mgr/>
- ▶ IBM Tivoli Data Warehouse:  
<http://www-3.ibm.com/software/tivoli/products/data-warehouse/>
- ▶ IBM Tivoli Enterprise Console:  
<http://www-3.ibm.com/software/tivoli/products/enterprise-console/>
- ▶ IBM Tivoli NetView:  
<http://www-3.ibm.com/software/tivoli/products/netview/>
- ▶ IBM Tivoli Business Systems Manager:  
<http://www-3.ibm.com/software/tivoli/products/bus-sys-mgr/>
- ▶ IBM Tivoli Service Level Advisor:  
<http://www-3.ibm.com/software/tivoli/products/service-level-advisor/>
- ▶ IBM Tivoli Data Warehouse:  
<http://www-3.ibm.com/software/tivoli/products/data-warehouse/>
- ▶ IBM Tivoli Enterprise Console:  
<http://www-3.ibm.com/software/tivoli/products/enterprise-console/>
- ▶ IBM Tivoli NetView:  
<http://www-3.ibm.com/software/tivoli/products/netview/>
- ▶ IBM Tivoli Business Systems Manager:  
<http://www-3.ibm.com/software/tivoli/products/bus-sys-mgr/>
- ▶ IBM Tivoli Service Level Advisor:  
<http://www-3.ibm.com/software/tivoli/products/service-level-advisor/>
- ▶ IBM Tivoli Monitoring:  
<http://www-3.ibm.com/software/tivoli/products/monitor/>

- ▶ IBM Tivoli Data Exchange:  
<http://www-3.ibm.com/software/tivoli/products/data-exchange/>
- ▶ IBM Tivoli Configuration Manager:  
<http://www-3.ibm.com/software/tivoli/products/config-mgr/>
- ▶ IBM Tivoli License Manager  
<http://www-3.ibm.com/software/tivoli/products/license-mgr/>
- ▶ IBM Tivoli Monitoring for Transaction Performance:  
<http://www-3.ibm.com/software/tivoli/products/monitor-transaction/>
- ▶ Integrated Identity Management:  
<http://www-3.ibm.com/software/tivoli/features/idmgmt/>
- ▶ IBM Tivoli Access Manager for e-business  
<http://www-3.ibm.com/software/tivoli/products/access-mgr-e-bus/>
- ▶ IBM Tivoli Directory Server:  
<http://www-3.ibm.com/software/tivoli/products/directory-server/>
- ▶ IBM Tivoli Directory Integrator:  
<http://www-3.ibm.com/software/tivoli/products/directory-integrator/>
- ▶ IBM Tivoli Identity Manager:  
<http://www-3.ibm.com/software/tivoli/products/identity-mgr/>
- ▶ Web Infrastructure Orchestration:  
<http://www-3.ibm.com/software/tivoli/features/web-svr-prov/>
- ▶ IBM Tivoli Provisioning Manager:  
<http://www-3.ibm.com/software/tivoli/products/prov-mgr/>
- ▶ How IBM Infrastructure Orchestration works:  
<http://www.ibm.com/news/us/2003/10/081b.html>
- ▶ IBM Tivoli Intelligent ThinkDynamic Orchestrator:  
<http://www-3.ibm.com/software/tivoli/products/intell-orch/>
- ▶ IBM Tivoli Monitoring:  
<http://www-3.ibm.com/software/tivoli/products/monitor/>
- ▶ IBM Director 4.1:  
[http://www-1.ibm.com/servers/eserver/xseries/systems\\_management/director\\_4.html](http://www-1.ibm.com/servers/eserver/xseries/systems_management/director_4.html)
- ▶ IBM Remote Deployment Manager:  
[http://www-1.ibm.com/servers/eserver/xseries/systems\\_management/sys\\_migration/rdm.html](http://www-1.ibm.com/servers/eserver/xseries/systems_management/sys_migration/rdm.html)

- ▶ IBM eServer BladeCenter:  
<http://www-1.ibm.com/servers/eserver/blades/>
- ▶ FASt T900 Storage Server:  
<http://www.storage.ibm.com/disk/fastt/fast900/index.html>
- ▶ Logical Partitioning for IBM iSeries:  
<http://www.ibm.com/servers/eserver/series/lpar/>
- ▶ iSeries Information Center - Logical partitioning:  
<http://publib.boulder.ibm.com/series/v5r2/ic2924/info/rzaj9/rzaj9ic1par.htm>
- ▶ Linux on iSeries:  
<http://www.ibm.com/servers/eserver/series/linux/>
- ▶ Linux on iSeries white paper:  
<http://www.ibm.com/servers/eserver/series/linux/pdfs/series.pdf>
- ▶ IBM eServer iSeries Capacity on demand:  
<http://www.ibm.com/servers/eserver/series/ondemand/cod/>
- ▶ IBM Tivoli System Automation for Linux:  
<http://www.ibm.com/software/tivoli/products/sys-auto-linux/>
- ▶ IBM Tivoli Systems Automation for OS/390:  
<http://www.ibm.com/software/tivoli/products/system-automation-390/index.html>
- ▶ Geographically Dispersed Parallel Sysplex:  
<http://www.ibm.com/servers/eserver/zseries/announce/april2002/gdps.html>

## How to get IBM Redbooks

You can search for, view, or download Redbooks, Redpapers, Hints and Tips, draft publications and Additional materials, as well as order hardcopy Redbooks or CD-ROMs, at this Web site:

[ibm.com/redbooks](http://ibm.com/redbooks)

## Help from IBM

IBM Support and downloads:

[ibm.com/support](http://ibm.com/support)

IBM Global Services:

[ibm.com/services](http://ibm.com/services)

# Index

## A

access 9  
    access control 69–71, 99–100, 180, 182, 184, 186, 188, 191–192, 198–199, 204–205, 209, 233  
    access control list 192, 195  
    access control management 22, 191, 198  
    access-control structures 200  
    block level access 36  
    Cisco Access Control Server 99  
    secure access 69–70, 99–100, 167–168, 170, 172, 176, 198, 204, 213  
    seeunder access  
    service access point 232  
    Tivoli Access Manager 72, 85, 100, 112, 176, 182, 184–185, 188, 190, 192, 197–200, 205–212, 220–221  
access control  
    see under access  
access control management 198  
ACL  
    see access control list  
    see under access control list  
Activity planner 129  
administrative interface 54  
Administrator's User Interface 152  
AIX 31–32, 49, 51, 54, 58, 60–61, 71, 79–80, 84, 89, 94, 147, 150, 198, 201, 203, 220–221, 244–245, 249, 263, 298, 303, 307, 310  
AME 83  
Apache 55, 220–221, 228, 297, 334  
Apache Web server 334  
application controller 226, 228, 235–238  
Application Extension Facility 129  
application provisioning 25  
Application Resource Management 75  
Application Response Method 151  
application server 205  
architecture 10  
ARM 54–55  
    see Automated Response Measurement  
Asset Manager 127  
asset relationship 233

auditing facilities 170  
authentication 21–22, 70, 99, 168–169, 176, 180, 184, 188–191, 197, 202, 205, 209–211  
authentication infrastructure 197  
authentication mechanisms 184  
authentication services 188  
authorization infrastructure 197  
authorization server 205  
authorization services 188  
automated attribute generation 200  
automated provisioning 226, 234  
Automated Response Measurement 54–55, 75, 151, 159–162, 259, 261, 265, 268, 297–298, 303–304, 307–308, 317  
automated security 167  
automation 7, 9, 12, 14, 16–19, 21–22, 24–25, 38, 59, 64–65, 76–78, 81, 83, 85, 88, 92, 96, 98–100, 109–110, 114, 118–119, 129, 139, 158, 160, 167, 172, 176, 216–218, 234, 240, 272, 278  
    elements 17  
autonomic 17–18, 20–21, 74, 77, 80, 83, 85, 88, 92, 98, 101, 108, 113–115, 125, 135, 144, 160, 162, 216, 224, 254, 321–322  
    autonomic computing 17–18, 74, 80, 83, 85, 101, 114–115, 144, 322  
    Autonomic Computing Toolkit 83, 85  
    Autonomic Management Engine 83  
autonomic computing 17–18  
Autonomic Computing Toolkit 83  
availability 9, 15–21, 32, 34, 36–37, 39, 58, 61–63, 78  
availability manager 113

## B

B2B 97  
B2C 97  
BEA WebLogic 211  
best practices 118, 149, 250  
BladeCenter 32, 51, 56–57, 61, 80, 140, 223, 228, 245, 252  
block level access 36  
Blueprint 101–102, 147, 160, 199  
blueprint

- identity management 199
- security 101
- BMC 57
- boot servers 231
- BPEL 126
- BPEL4WS 336
- business
  - business flexibility 6–8
  - business goal 7, 98, 259, 299, 305
  - business modeling 8, 345
  - business objectives 9, 15, 79, 107, 125, 157, 159, 175, 216, 221, 276–277
  - business process 7–10, 25, 69, 75, 77–78, 88, 108, 117–118, 126, 159, 168, 336
  - business service management 9, 16, 18, 109–110, 240
  - business strategies 14
- business drivers 14
- business flexibility 6–8
- Business Impact Management 111
- business policies 19
- business process 7, 9–10, 25
- business process management 9
- business processes 125
- business service management 9, 16, 18
- business strategies 14

## C

- CA 57
- cache 192
- capabilities 7
- Capacity BackUp 89, 241–244
- Capacity BackUp on Demand 241
- Capacity on Demand 24, 32, 50, 89, 241–242, 244–246
- Capacity Upgrade on Demand 24, 89, 241, 244
- CAPP 71
- CBE
  - see common base event
- CBU
  - see Capacity BackUP
- certificates 184
- CICS 48, 92
- CIFS 63
- CIM 38, 116, 148, 337, 339, 346–347
- Cisco 21, 23, 39, 91, 98–100, 102, 119, 256–257, 259–260, 263, 267
- Cisco Access Control Server 99

- Cisco Network Admission Control (NAC) 99
- Cisco Security Agent 99
- cluster 24–25, 31–32, 57, 79, 116, 120, 146, 217, 225–227, 230, 232, 235, 238–240, 256, 269, 279, 282, 287–288, 290–291
- cluster manager 225
- cluster organizations 232
- cluster visualizing 226, 232
- Clustering 27, 31–32, 48, 89
- CMM 148
- collaboration 9, 98, 191
- common base event 17, 74
- Common Information Model 116, 127, 337, 346
- Common Internet File System 63
- Common Management Model 148
- Component Business Model 4
- configuration changes 113
- Connectivity on Demand 241
- consistent management 73
- consolidation 15, 29–31, 35, 59, 71, 103, 131–140, 149, 152–153, 166, 243, 245, 249, 271–272, 277–279, 282, 293, 330
  - full consolidation 135
  - physical consolidation 15, 35, 132–133, 136, 138, 152
  - physical storage consolidation 133
  - virtual consolidation 135
- Content Switching Module 255–256
- Control Center 53, 78, 259, 264–265, 269, 296, 301–302, 304, 308, 310–311, 313, 317
- control operations 74
- CORBA 338
- correlation server 103
- CSM 256–263, 265–269
- CUoD 50, 89, 241–242
- CuOD 50, 89, 241–242
- custom realm 211

## D

- Daily On/Off Capacity on Demand 24
- DAS 35, 120
- dashboard 111
- data acquisition engine 226
- Data Center Model 227, 232, 287–291
- data consolidation 132
- data management 146
- data protection 170, 172
- DB2 39–40, 55, 75, 83, 91–92, 95, 101, 108, 118,

- 121, 183, 201, 204, 220–221, 223, 228, 230, 263–264, 297–298, 304–306, 308–309, 318
- decision making 226, 232
- delegated administration 200
- delegation of authority 170
- demilitarized zone 174
- denial of service 71
- denial of service attacks 97, 171
- deployment engine 227
- deployment process 118
- derived data 232
- DFSMS 23
- DFSMSHsm 23
- Direct Attached Storage 35, 120
- direct attached storage 120
- Director Multiplatform 56–57, 64–65, 116, 127, 130, 243, 327
- directory integrator 205
- directory server 201, 205
- directory server replica 189
- distributed systems 16
- DLPAR 135–136
- DMTF 334, 337
- DMZ 185, 193, 235
- domain manager 54–56, 151–152, 228, 258–259, 261–265, 267–269, 297, 300–301, 303–304
- domain policy 54, 56, 151–152, 259, 298, 301–302, 304–306, 308–313
- Domino 55
- drivers 252
- DS4000 35, 82, 134, 138, 140, 280, 290
- DS8000 35, 134, 138, 140
- dynamic re-provisioning 215
- dynamic sense and respond 21

## E

- Eclipse 334
- ED/FI 93
- electronic service agent 82
- elements of automation 17
- encryption 71
- endpoint support 200
- end-to-end 5, 9–10, 14, 20–22, 27, 33, 42, 52, 54, 61
- end-to-end visibility 74
- Enterprise Identity Mapping 202
- Enterprise Service Bus 11, 88, 108
- Enterprise service bus 11, 88, 108

- Enterprise Workload Management 23, 150–151
- Enterprise Workload Management Control Center 78, 259, 264–265, 269, 301, 304, 308, 310–311, 313, 317
- Enterprise Workload Management domain 259, 261, 264, 267, 300
- Enterprise Workload Manager 46, 52–53, 64, 75, 77, 85, 91, 117, 119, 123, 145, 150–151, 155, 161–162, 243, 255–256, 268
- Enterprise Workload Manager Control Center 53
- ERP 201
- Error Detection and Fault isolation 93
- ESB
  - see Enterprise Service Bus
- esb 11
- ESS 35, 82, 122, 138, 140, 246, 280–281
- ESX 146, 247
- Ethernet 32, 37, 47, 50, 89, 244, 249, 274, 287
- event management 21

## F

- FAStT 35, 138, 140
- FAStT900 Storage Server 223
- FC 242, 281, 287
- Federal Information Processing Standard 71
- Fiber Channel 138
- file level access 35
- file movement 122, 139, 154, 271, 278, 290
- Fileset 282–283, 285, 288, 290, 293
- FIM 72, 212
- FIPS 71
- framework 1, 8, 10, 13–14, 16–19, 21, 24–25, 32, 42, 55, 57, 67, 78, 147–148, 200, 213, 216, 254, 272, 293, 343, 346
  - infrastructure management 14

## G

- GAR 147
- GDPS 20, 90, 204
- Generic Log Adapter 83
- GGF 334, 337
- Global Grid Forum 144, 334, 337, 343
- global grid forum 337
- global resource manager 226
- Globus 42, 60, 120, 146–147, 149
- Globus Alliance 42, 60
- Globus Toolkit 42, 120, 147
- Globus toolkit 120

- graphic interface 117, 275
- grid 16, 25, 32, 39, 41–42, 60, 76–77, 85, 115, 120, 123, 135, 144, 146–149, 155, 218, 247, 327, 334, 337–338, 343–345
- Grid Computing 135
- grid computing 16, 144, 345
- Grid Core Services 144
- Grid Data Services 144
- Grid Packaging Technology 148
- Grid Program Execution Services 144
- Grid Service Archive 147
- Grid services 77
- Grid Solutions 115
- Grid Toolbox 85, 247

## H

- HACMP 21, 32
- health center 79, 90
- Health Console 95, 162
- health console 95
- heterogeneous 14, 16, 18–21, 23, 26–27, 32, 36, 41–43, 45–46, 52–53, 56–57, 60–61, 63, 65
- HiperSockets 26, 38, 48, 90, 120, 149, 248
- HP OpenView 57
- HP-UX 94, 201
- HTTP 30, 39, 54–55, 70–71, 147, 151, 168, 191, 207, 210, 220–221, 228, 230, 238–239, 263, 296, 334, 341, 347
- HyperSwap 90
- hypervisor 30, 47, 49–50, 76, 135–137, 158, 244–248, 253, 322

## I

- i5 29, 49–51, 58, 60, 245
- i5/OS 49, 51, 58, 60, 245
- IBM Cloudscape 147
- IBM Cluster System Management 31
- IBM Director 39, 51, 56–57, 64–65, 79–80, 85, 89–90, 96, 116, 123, 127, 130, 223, 228–230, 240, 243, 252, 327
- IBM Director Multiplatform 56–57, 64–65, 127, 130, 243, 327
- IBM Directory Integrator 208
- IBM Directory Server 201
- IBM Dynamic Infrastructure Enterprise Edition for mySAP Business Suite 55, 60–61
- IBM Enterprise Workload Manager 52, 64, 85, 117, 119, 155, 255

- IBM Globus Toolkit 42
- IBM Grid Toolbox 60, 85, 123, 146–147, 155, 247
- IBM Tivoli Comprehensive Network Address Translator 80, 85
- IBM Tivoli Configuration Manager 127, 129–130
- IBM Tivoli Enterprise Console 80, 85, 92, 96, 102, 105, 110–111
- IBM Tivoli Intelligent ThinkDynamic Orchestration 59
- IBM Tivoli Monitoring 57, 79, 81, 85, 91–92, 94–96, 110–112, 116, 161–162, 223, 230
- IBM Tivoli Provisioning Manager 39, 58–59, 64, 99–100, 105, 117–118, 123, 149–150, 223, 250, 252, 327
- IBM Tivoli System Automation 81, 83, 85, 92, 96
- IBM TotalStorage DS6000 35, 140
- IBM TotalStorage Enterprise Storage Server 35, 82, 140
- IBM TotalStorage File System 43, 122–123, 139–140, 153, 274, 278, 282, 284–291, 293
- IBM TotalStorage Productivity Center 61–62, 64, 81–82, 85, 93, 96, 120–121, 123, 139–140, 274, 277–278, 285–286, 291, 293
- IBM TotalStorage Productivity Center for Data 82, 93, 96, 120–121, 123, 139, 278, 285–286
- IBM TotalStorage Productivity Center for Disk 82, 140
- IBM TotalStorage Productivity Center for Fabric 82, 93, 96, 120–121, 123, 139, 278, 285
- IBM TotalStorage Productivity Center for Replication 82, 140
- IBM TotalStorage SAN File System 37, 63–64, 278, 282, 293
- IBM TotalStorage SAN Volume Controller 43, 62, 64, 121–123, 138, 140, 153, 155, 274, 277–282, 284, 286–292
- IBM TotalStorage Volume Controller 36
- IBM Virtualization Engine 1, 23, 43, 45–47, 52, 58, 60–61, 63–65, 79, 84–85, 96, 110, 112, 123, 163, 243, 296–297, 302, 327, 331–332
- IBM Virtualization Engine Suite for Servers 43, 52, 60, 64, 79, 84–85, 297
- IBM Virtualization Engine Suite for Storage 43, 52, 61, 64
- IDE 35, 334
- identity control layer 200
- identity management 22, 72, 170, 181–183, 185–186, 197, 199–202, 212, 219
- identity management blueprint 199



- identity mapping 25
- identity synchronization 183, 205
- IETF 334
- IGT3 146–147, 149
- IIS 55, 297–298, 308–309
- IMS 92
- industry standards 5
- information gathering 226, 230, 287–291
- Informix 95
- infrastructure management 1, 7, 9–10, 12–14, 16, 19, 25, 34, 43, 61, 67, 81, 114, 116, 128, 132, 134, 139, 148–149, 159, 165–166, 217, 222, 240, 251, 272, 276, 285, 291, 323
  - framework 14
  - overview 13
- instrumentation 38–39, 55, 148, 259, 297–298
- Integrated Drive Electronics 35
- Integrated Facility for Linux 120
- integrated identity management 181
- integration 6, 8, 10–12, 14, 40–42, 57, 64
- integration capabilities 8
- Intel 26, 32, 51, 57, 119–120, 140, 201
- Intelligent Resource Director 24, 33, 48, 90, 146, 155, 241
- Internet SCSI 36
- introduction 3
- intrusion detection 98, 101
- intrusion prevention 200
- intrusions 97
- IRD 33, 48
  - see Intelligent Resource Director
- ISC 75
- iSCSI 36, 138–140
- iSeries 26, 29, 32, 50, 55–61, 83, 89, 92, 120, 136, 140, 146–147, 201, 203, 242, 244–245, 247, 249
- IT simplification 7–8, 16, 25
- ITITO 59

## J

- J2EE 75, 161, 210–211
- J2EE application deployment 210
- Java 48, 95, 147–148, 162, 201, 269, 334–336, 338
- JCA 148
- JCP 334–335
- JNDI 208
- job scheduling 218
- JS20 140

- JSR168 75
- junctions 189

## K

- key standards 338

## L

- LDAP 169, 183
- Life Cycle Management 39, 99, 122, 134, 154, 271, 276–278, 284–285, 290, 293
- lifecycle management 199
- Lightweight Third-Party Authentication 209
- Linux 29–31, 48–49, 51, 54, 59–61, 63, 79–80, 83–84, 92, 94, 96, 118–120, 123, 137, 139–140, 147, 149–150, 198, 201, 203, 220–221, 242–245, 248–249, 263, 274, 288
- load balancer 225, 231, 238–239, 256–260, 263
- load balancing 218
- Load Balancing Advisor 256
- Log and Trace Analyzer 83
- logical consolidation 132–134, 138–139, 152–153, 277–279, 282
- logical partitioning 27, 29–31, 46–48
- Logical Unit Number
  - see LUN
- Lotus 55
- LPAR 24, 27, 29–31, 46–48, 50, 59, 89, 123, 135–136, 145–146, 149, 242–243, 245
- LPAR weights 145
- LPTA 209
- LUN 20, 35, 121, 134, 279–281, 288–289

## M

- managed security services 104
- managed server component 151
- management interface 227
- manager server 54, 56, 151–152, 258–259, 262, 267–269, 301, 303–304, 311, 313–314, 317–320
- managing capacity 224
- managing risk 98
- MDA 338–339, 346
- metadirectory connectivity 200
- microcode 8, 29
- Micro-Partitioning 29, 49
- micro-partitioning 29, 49, 89, 244
- Microsoft 30, 37, 51, 55, 57–58, 61, 148, 184, 201, 243, 249, 298, 342–343

monitor systems 107, 159  
mySAP.com 95

## N

NAS 35, 120, 138–139, 278  
NetBEUI 37  
NetIQ 57  
Network Attached Storage 35, 120, 138  
network attached storage 120, 172  
network dispatcher 189  
network utilization 218  
network virtualization 16  
NIM 49  
Nortel 23, 39, 119  
Novel IPX 37  
N-tier applications 157, 216

## O

OASIS 32, 74, 334, 336–338, 342–343  
OGSA 41, 60, 76–77, 115, 120, 126, 144, 146–148, 337–338, 343–344  
OGSI 76, 147, 337–338, 344  
OMG 334, 338  
OPAL 59  
Open Grid Services Architecture 41, 60, 76, 115, 144, 337, 343  
Open Group 54, 151, 297, 347  
open source 334  
open standards 169, 217, 273, 297  
Operating Environment architecture 10  
operating systems 8, 24, 29–30, 34, 36, 38–39, 45–46, 58–59  
optimization 9, 16, 22  
optimize utilization 143  
Oracle 95, 121, 128–129  
orchestration 7, 9, 16–17, 19–20, 59, 76, 78, 109, 115, 144–145, 152, 157, 216, 222–224, 226, 229–230, 240, 251–252  
Orchestration and Provisioning Automatic Library 59  
orchestrator 19, 23, 76, 81, 85, 93, 96, 117, 123, 150, 152, 160, 223–224, 226–227, 229–232, 234, 240, 250–252  
OS/390 29, 92, 96  
OS/400 54, 71, 79, 84, 94, 120, 147, 249

## P

p5 29, 244  
Parallel Sysplex 20, 31, 48, 84, 90, 92, 96, 149, 204  
partitioning 27–28, 46  
partitioning 26–32, 38, 46–50, 89, 120, 135–136, 146, 149, 242–245  
    logical partitioning 26–27, 29–31, 46–48, 50, 135–136, 146, 242–243  
    physical partitioning 27–28  
password change/reset 200  
performance measurement 152  
persistent universal auditing 200  
physical consolidation 15, 35  
physical partitioning 27–28  
PKI 169  
PLM 33, 49  
plug-in 51, 55, 63, 297–298, 308–309  
policy 9, 11, 16–17, 19–22, 24, 38–39, 54, 56, 61, 63–64, 70, 77–79, 81–82, 87–88, 91, 93, 99–100, 104, 108–109, 117, 122, 134, 136, 139, 148, 151–152, 154, 160, 162, 168–170, 175, 182–186, 188, 190, 192–195, 198, 200, 210–211, 218, 238, 247, 251, 259, 271, 278, 282–285, 288–289, 293, 298–302, 304–306, 308–313, 315  
    policy service agent 148  
    policy service manager 148  
policy based orchestration  
    orchestration 19, 109  
policy server 185  
policy-based orchestration 9  
pool resources 143  
POWER4 243  
power5 29, 49–50, 245  
predictive management 111  
privacy control 170–171, 182  
privacy control management 203  
privacy management 22  
Proactive Analysis Components 92, 95  
problem determination 17, 77, 80, 83, 88, 102, 108, 162  
problem isolation 152  
process class 307, 310–311, 313–314  
process transformation 8  
protect systems 97  
provision system resources 157, 215, 271  
provisioning 7, 9, 16–17, 19–21, 24–25, 33, 38–39, 46, 58–62, 64, 70, 76–78, 81, 88, 91, 98–100, 105, 109, 115, 117–119, 123, 132, 135, 144–145, 149–150, 152, 157–158, 160, 162, 166, 168,

177–178, 182–183, 193, 198, 201, 205–208,  
215–218, 223–227, 230, 232, 234, 238–241,  
250–254, 327  
provisioning infrastructure 224  
provisioning scenarios 234  
pSeries 26, 29, 31–33, 49–50, 56, 58–61, 71, 83,  
89, 92, 119–120, 136, 140, 146–147, 150, 201, 203,  
243–244, 247

## Q

Quality of Service 19, 37  
quality of service 11–12, 14, 16, 19, 26, 37–38, 40  
query tool 127, 130

## R

RAID 89, 133, 140, 153, 280, 282  
real time management 110  
Redbooks Web site 354  
    Contact us xx  
Remote Deployment Manager 223, 229  
replicas 225  
replication of services 188  
reporting 200  
resiliency 17, 20, 87–88, 140, 250  
resilient 3, 10, 60, 167, 169, 215, 217, 273, 296  
resource management 146, 222  
resource monitors 239  
resource pool 24, 117, 232, 235, 237–238  
resource pools 24, 232  
resource virtualization 9  
revoking existing user 196  
risk analysis 146  
risk management 192  
RMF 84–85  
Roadmap 70, 168  
role-based access-control structures 200  
roles 170  
root-cause analysis 21  
RS/6000 31  
rule sets 170  
runtime environment 17, 147–148, 303

## S

SAML 72, 212  
SAN 34, 36–37, 62–64, 82, 93, 120–122, 128, 134,  
138–140, 145, 152–155, 172–174, 219, 224–225,  
231, 241, 274, 276–279, 282, 285, 287–289,

291–293  
SAN Error Predictor 93  
SAP 48, 60–61, 83, 92  
scalable 50, 63, 73, 169, 171, 187–189, 199, 217,  
221, 273, 296, 335  
scale out 32  
SCSI 35–36, 50, 134, 138, 243  
SDK 147  
secure access 69, 167  
Secure Socket Layer 70, 168–169  
security 9–10, 16–17, 19, 21–22, 25, 30, 37, 39–40,  
42, 61, 70–72, 77, 97–105, 108–109, 135, 137,  
146–147, 149, 166–172, 174–176, 178–180,  
183–184, 186–195, 197–198, 200–201, 203–204,  
210, 212, 220–221, 227, 240, 251, 267, 288, 296,  
335–339, 343  
security architecture 179  
Security Assertion Markup Language 72, 212  
security benefits 186  
security management 97  
security management console 170–171  
Security On Demand Infrastructure 186  
security policies 188  
self-care 196, 200  
self-healing 129  
self-managing 9, 17, 100, 108, 169, 216, 254, 273,  
296, 322  
self-tuning 53  
Server Allocation for WebSphere Application Server  
252  
server virtualization 16  
Server/Application State Protocol 256, 258,  
260–262, 265, 269  
service access point 232  
service classes 53–54, 79, 91, 162, 256, 305–320  
service level 15, 17–19, 25, 38, 54, 56, 88, 91, 103,  
107–109, 111–112, 125, 143, 150–152, 158, 160,  
216, 222–224, 227, 230, 235, 238, 296, 298–299,  
301, 305–306, 321, 331  
Service Level Agreements 19  
service level management 108  
service levels 222  
Service Oriented Architecture 11, 76, 115  
Service oriented architecture 7, 11–12, 76  
service policy 19–20, 24, 310–311, 315  
SGML 336  
Siebel 95  
Siebel eBusiness Applications 210  
Simple Object Access Protocol 41, 335, 341

- simplification 7
- simplify 131
- single sign on 72
- single sign-on 72, 184, 198, 209
- SLA 19, 24, 56, 61, 84, 88, 107–108, 111, 152, 222, 228, 298–299
- Small Computer System Interface 35
- SNA 37, 84
- SNMP 38, 80, 82, 111, 119, 148, 232, 289
- SOA 11, 76, 115
- SOA (see Service Oriented Architecture)
- SOAP 41, 60, 70, 148, 168, 335, 338–339, 341–343
- software package blocks 129
- Software Package Editor 130
- software partitioning 27, 30
- Software Repository 127
- software stacks 232
- Solaris 54, 58, 94, 201, 263, 298, 303
- SQL 121, 162, 282
- SSL 70–71, 168–169, 185–186
- standard 12, 17–18, 29, 32, 36–37, 43, 45, 51, 54, 59, 71, 74–75, 77, 122, 138, 144, 146, 148, 151, 159–160, 192, 201, 212, 245, 274, 282, 297, 334, 336–337, 340–342, 344–346
- standards organizations 334
- standards overview 333
- Standby Capacity on Demand 241, 246
- Stocks-4u.com scenario 173, 218
- Storage Area Network 34–37, 82, 85, 93, 96, 120, 139
- Storage Infrastructure Management 81
- storage management 81
- Storage Pool 62, 121–122, 154, 277, 280, 282–284, 288–290, 293
- storage usage 23
- storage virtualization 16, 138, 277
- Subsystem Device Driver 288
- SUN 55, 58, 298, 335
- Sun 55, 58
- System Automation (see Tivoli System Automation)
- system failures 87
- system management server 229
- system resource failure 240
- system utilization 158, 216
- systems services 43, 45–47, 52, 63, 332
- systems technologies 43, 45–47, 49–51

## T

- TAI 209
- TCG 212
- TCO 15, 61, 132, 329–330
- TCP/IP 29, 37, 48, 84, 104, 138, 204, 261
- TCPA 212
- TEC 82
- terminal servers 231
- The Trusted Computer Platform Alliance 212
- ThinkVantage Technology 99
- Tivoli Access Manager 182, 184, 199, 205, 207, 209
- API 192
- Tivoli Business Service Management 109
- Tivoli Business Systems Manager 79, 92, 103
- Tivoli Comprehensive Network Address Translator 80
- Tivoli Data Warehouse 103
- Tivoli Directory Integrator 183
- Tivoli Directory Server 183
- Tivoli Enterprise 57, 80, 82, 85, 92, 95–96, 102–103, 105, 110–111
- Tivoli Enterprise Console 80, 82, 85, 92, 95–96, 102–103, 105, 110–111
- Tivoli Identity Manager 99, 183, 194, 196, 199, 205, 207–208
- Tivoli Integrated Identity Management 182
- Tivoli Intelligent ThinkDynamic Orchestrator 81, 117, 150, 250
- Tivoli Intrusion Manager 101, 103
- Tivoli Management Region 129, 230, 239
- Tivoli Monitoring 81, 91, 229
- Tivoli NetView 57, 102, 110–111
- Tivoli Privacy Manager 182
- Tivoli Provisioning Manager 118, 149, 250
- Tivoli Risk Manager 80, 98, 101–103
- Tivoli SAN Manager 93
- Tivoli Security Event Management 101, 104
- Tivoli Security Management blueprint 101
- Tivoli Service Level Advisor 103, 111–112, 160
- Tivoli Storage Manager 252
- Tivoli Storage Resource Manager 121
- Tivoli System Automation 77, 83, 92
- Tivoli Systems Automation 21, 96
- Tivoli Web Site Analyzer 111
- Tivoli Workload Scheduler 92
- TLS 169
- TMTF 91, 110, 160–163
- Token Ring 37

- toolkit 42, 60, 83, 85, 120, 147
- topology 18, 21, 35, 54, 75, 78, 82, 91, 102, 139, 151, 159, 161, 258, 262–264, 267, 278, 285, 291, 296–297, 302, 311–312, 314, 317–318, 320–321, 332
- total cost of ownership 132
- TotalStorage 241, 246
- transaction class 307, 309, 311, 313–314, 316
- transaction management 20
- trust association interceptor 209
- Trusted Computer Platform Alliance 212
- Trusted Computing Group 212

## U

- UDDI 336, 338, 342
- UML 338, 345
- uniform security model 180
- UNIX 28, 30, 37, 48, 63, 139, 153, 184, 198, 204, 243, 274, 288
- US Government Orange Book B1/C2 71
- user provisioning 182–183, 198, 205, 208
- user self-registration 196

## V

- VE 43, 47, 61, 65, 89, 116, 145, 150, 303, 325–331
- VE Console 89, 116
- Virtual Ethernet 47, 50, 89, 244, 249
- Virtual I/O 47, 49–50, 89, 244, 249
- virtual IP address 218
- Virtual Machine Manager 51, 89
- Virtual SCSI 50
- virtualization 7, 9–10, 12, 14, 16, 23, 25–40, 42–43, 45–47, 49, 51–52, 56–58, 60–61, 63–66, 215
- Virtualization Engine 1, 23, 42–43, 45–47, 50, 52, 56–58, 60–61, 63–65, 75, 79, 84–85, 90, 96, 110, 112, 123, 163, 243, 263, 296–297, 302, 325, 327, 329, 331–332
- Virtualization Engine Console 57–58, 64, 75, 79, 85, 90, 96, 123, 327
- virus threats 97
- VLAN 38, 49–50, 239, 248
- VMWare 30, 51, 59, 89, 119, 136, 140, 146, 247
- VMware 136, 140, 247
- VPN 38, 71, 102
- VSWITCH 248

## W

- W3C 334–335
- WAS 5–6, 23, 29, 31, 35–36, 43, 85, 112, 136, 145–148, 150, 153–154, 166, 173, 176, 204–205, 218, 256–257, 280, 290, 298, 302, 318, 321, 334–336, 338–339, 342
- WBEM 337, 339
- Web Infrastructure Orchestration 222, 252
- Web Service Resource Framework 148
- Web Services 12, 32, 39–42, 69–70, 72, 88, 144, 147–148, 168–169, 183, 197–198, 201, 212, 334–339, 341–345
- Web Services Description Language 41, 147, 335, 342
- Web Services Resource Framework 42
- Web Services Security 70, 72, 168–169, 212, 336, 343
- Web site intelligence 23
- WebSEAL 176, 185, 193, 205, 209, 235
- WebSphere 108
- WebSphere Application Server 118, 205, 238
- WebSphere Application Server 19, 23, 25, 55, 75, 91, 118, 189, 205, 209–210, 220–221, 223, 228, 230, 238–239, 252, 263, 298, 301, 303–304, 308
- WebSphere Application Server Network Dispatcher 230
- WebSphere Edge Server 162, 189
- WebSphere Interchange Server 95
- WebSphere MQ 95
- WebSphere MQ Workflow 206
- weight 29, 145, 256–259, 261–262, 265–269
- Windows 28, 30–31, 51, 54, 58–59, 61, 63, 79–80, 82, 84, 94, 99, 139, 150, 184, 198, 201, 243, 245, 249, 263, 265, 274, 288, 298, 310
- Windows 2000 94, 201
- Windows NT 94
- WLM 19–20, 23–25, 31, 33, 48, 77, 84, 90, 145, 150, 155, 162, 256
- WLM (see Workload Manager)
- workflow 118, 200, 206, 232, 234
- workload management 19, 77, 136, 215
- Workload Manager 33, 46, 48, 52–53, 64, 75, 77, 84–85, 90–91, 113, 117, 119, 123, 136–137, 145–146, 150–151, 155, 158, 161–162, 217, 243, 247–248, 253, 255–256, 260, 268, 321–322
- workload manager 113, 246
- WS-\* standards 12
- WSDL 41, 60, 147–148, 335, 338–342, 344
- WS-I 334, 336

WS-I Basic Profile 337–338, 342  
WS-Policy 70, 168  
WS-RF 42, 148  
WS-SecureConversation 168  
WS-Security 70, 168, 336–339, 343

## **X**

Xalan project 334  
Xeon 140  
XHTML 335  
XML 12, 40–41, 54, 60, 74, 148, 151, 234,  
334–336, 338–342, 344–345, 347  
XML Schema 341  
XML standards 339  
xSeries 30–31, 38, 51, 56–61, 83, 89, 92, 99,  
119–120, 130, 136, 140, 146–147, 150, 203, 223,  
228, 245, 247  
XSLT 334, 341

## **Z**

z/OS 19–20, 23, 29–31, 48, 54–55, 71, 84–85, 90,  
92, 104–105, 112, 120, 145–146, 149, 155, 162,  
198, 203–204, 220–221, 256, 321  
z/VM 30, 48, 59, 90, 120, 123, 136–137, 140,  
149–150, 243, 248–249  
zSeries 19–20, 23–26, 29–33, 38, 47–48, 56,  
59–61, 71, 77, 83, 89–90, 92, 120, 123, 136–137,  
140, 145–147, 149, 151, 155, 201, 203–204,  
242–243, 247–248



# On Demand Operating Environment: Managing the Infrastructure (Virtualization Engine Update)

(0.5" spine)  
0.475" <-> 0.875"  
250 <-> 459 pages









# On Demand Operating Environment: Managing the Infrastructure (Virtualization Engine Update)



**Redbooks**

## **Introduction to the On Demand Operating Environment**

This IBM Redbook (along with its companion volume, *On Demand Operating Environment: Creating Business Flexibility*, SG24-6633), provides an insight into the kind of operating environment required to support an On Demand Business.

## **Automate and virtualize the IT environment**

It provides an overview of the architecture of an On Demand Operating Environment and describes in more detail the components that are required to manage the infrastructure. To meet the business needs of being responsive, variable, focused, and resilient, an On Demand Operating Environment must be integrated, autonomic, virtualized, and open. Though these attributes are all interrelated, this redbook focuses on the automation and virtualization components as they enable efficient infrastructure management.

## **Start deploying today**

This redbook provides descriptions of several approaches that one can choose to start implementing pieces of an On Demand Operating Environment today. Which approach is right for the reader will depend on their specific business environment and their immediate needs.

## **INTERNATIONAL TECHNICAL SUPPORT ORGANIZATION**

## **BUILDING TECHNICAL INFORMATION BASED ON PRACTICAL EXPERIENCE**

IBM Redbooks are developed by the IBM International Technical Support Organization. Experts from IBM, Customers and Partners from around the world create timely technical information based on realistic scenarios. Specific recommendations are provided to help you implement IT solutions more effectively in your environment.

**For more information:**  
[ibm.com/redbooks](http://ibm.com/redbooks)